

Non-Additive Approaches to Aggregation

Ph.D.

Philosophy Department

Abstract

Sometimes we ought to aggregate lesser harms to many such that they outweigh greater harms to a few, and sometimes we ought not to. This seems self-evident, but it has proven surprisingly difficult to construct a coherent moral theory out of this basic observation. In particular, it is difficult to explain (in a principled way) *when* we ought to aggregate. Relevance views attempt to solve this problem by arguing that sufficiently lesser harms are irrelevant to greater harms and thus should not be taken into consideration. My thesis develops and defends a new relevance view.

The first half of the thesis focuses on when lesser harms can aggregate and *outweigh* stronger harms. I start by solving a ‘fatal’ dilemma for all relevance views. I show that one horn of the dilemma is fatal but avoidable, whilst the other is unavoidable but not fatal, and that the dilemma rests on a faulty assumption of weak-additivity. I then develop a new relevance view based around the principle that we ought to balance harms by similarity (i.e., we ought to balance greater harms against similarly great harms and lesser harms against other lesser harms) before deciding whether lesser harms are relevant to greater harms.

The second half focuses on when lesser harms can aggregate and *tiebreak* between stronger harms. First, I show that relevance views cannot capture our intuition that sufficiently lesser harms should not tiebreak between greater harms, without forfeiting our intuitions in other important cases. Secondly, I demonstrate that relevance views are sometimes inappropriately insensitive to the number of harms in tiebreaking cases. To solve these problems, I distinguish between two types of relevance, and I introduce Broomean fairness into my account. Doing so allows my view to capture all our intuitions about when lesser harms can tiebreak between stronger harms.

Keywords: Aggregation, Relevance, Fairness, Lotteries

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

James Hart

March 2023

For Philip

Acknowledgements

Very rarely in life do you get an opportunity to properly thank people, especially not in such a public way. I have many people to thank for help in writing my thesis, but I also want to take this opportunity to thank all the others who have supported me in wider ways. I owe a debt to all of you.

Firstly, my family. Mum, for being my first (and still best) teacher. You grew my curiosity and made me the person I am today (for better or worse). I'm sorry for all the tantrums I had when you made me write - I think I've got over that now. I promise to get a real job soon. Dad, for being the best role-model I could ever have, for our long debates that inspired a love of philosophy in me, and for reminding me to remain silly. My brothers, Peter, Jonny, Phillip and Ben, my best friends and a constant source of joy in my life. Chatting nonsense with you is my favourite thing in the world. Grandma, for taking an interest and being proud of me despite having no idea what I do, and for the financial support that helped make this possible. Granny B, for your constant prayers and invaluable spiritual wisdom. And the wider family, all who have supported and loved me in such myriad ways.

Secondly, my supervisors. I was lucky to have three amazing supervisors. All such incredible minds, deserving of great egos but never accompanied by them. Your comments have been tremendously helpful. Thanks for sticking with the project even though I pivoted halfway through to working on an area slightly outside your expertise. Your wide-ranging guidance and support going far beyond the PhD has made everything that comes next possible. Brad, thank you for staying on despite retiring and Jonathan for joining the project later but taking to it so wholeheartedly. Philip, I only wish you were here so I could thank you myself for being the best supervisor I could have ever hoped for.

The wider department at Reading also deserves more than a mention. You have been my home for close to 10 years now, and I have not once had a bad experience with the department. In my time here I have had an opportunity to be taught by or work with each and every one of you. The opportunities you have provided, the seriousness with which you treated me even as an undergraduate, and the support each step of the way has been very precious to me. I shall miss bumping into each of you on the corridor.

Fourth, to all the other PhD students at Reading, Southampton, and the wider DTP, for the comments, discussions, and solidarity. Special thanks to Steven Wu, my accountability buddy and now good friend. Our long discussions have been hugely helpful, not just in developing ideas but in sometimes reminding me why I do philosophy in the first place. I shall not thank you though for making me feel perpetually underread.

Fifth, my research was generously supported by the Arts and Humanities Research Council, via a studentship awarded to me by the South, West and Wales Doctoral Training Partnership [Grant Number: AH/R012776/1]. I owe the DTP thanks for believing my project to be worth funding, for all the opportunities they have made possible, and the relationships that would never have been if it weren't for them bringing us together. Similarly, Effective Thesis for coaching, support, and community. If any impact comes out of this thesis, it will be largely due to their guidance.

Many people looked at draft chapters and papers of the thesis and provided helpful comments despite receiving nothing for it. In particular, John Broome for happily taking the time to look through a critique of his work sent out of the blue, and Patrick Tomlin for comments on some of the early ideas long before this became the topic of my thesis. I should also thank Patrick for seeing something in me as an undergraduate, taking me to the Oxford Moral Philosophy Seminar (where I proceeded to spill wine over the visiting speaker) and for encouraging me to pursue this career. A number of anonymous reviewers have also provided some of the most powerful critiques of my work and forced me to make my work much better.

Seventh, my wider friends in Reading and beyond. You have all provided great distractions from the PhD and moral support whenever I have needed it. Life would be dull without you.

Beth, I owe you more than most on this list combined. Thank you for putting up with me, for humouring me about philosophy, for taking the risk on me and the PhD, for your misplaced confidence in my ability to make something out of this, for making it financially possible, for the help with the day-to-day grind, for rarely getting annoyed by me, and for the long games of Wingspan. I could never deserve or repay the love you give me.

Lastly, I owe everything to God. You are my Rock, my Guide, my ultimate Grounding and Purpose. Your blessings are overwhelming, and I thank you for each and everyone above.

Dedication

There is something peculiar about a supervisory relationship. The relationship is primarily, at least initially, a professional one. Your supervisor is an authority figure and in some sense your boss. Yet at the same time a PhD is such a personal, creative, and all-encompassing endeavour, that the project ends up meaning so much more. By the end of the project, you are not just putting your work forward for examination but you whole self. Thus, the PhD project does not remain entirely a professional one. The comments, criticism, and support you receive from your supervisor shape not just your work but who you are. As such, your relationship with your supervisor also shifts. They come to know you in a way that nobody else does.

I don't think I would have had cause to reflect on how my relationship with my supervisors had shifted until Philip's death. Philip was so kind, so invested, and so generous with his time. He had such a fierce intellect accompanied by a humility and willingness to engage with you, which meant he was never intimidating in the way one might expect of such a brilliant mind.

I would often see his office door slightly ajar and knock hoping simply to say hello, let him know some small development or ask for a minor piece of advice or admin. We would invariably talk for about five minutes in the doorway before he asked me to take a seat. The armchair in his office sank with the weight of students that had been there before and reflected how comfortable Philip made his students. The quick hello would often become hours spent chatting in that office. Philip was immensely generous with his time and never made you feel like an imposition.

Of course, we spoke philosophy and my project, but we also spoke about everything else. He would listen to my fears and ambitions, take interest in the bland parts of my life and he treated my achievements as if they were his own. This was not a one-way relationship; he would speak of his life and his family, his excitement for Alec's wedding, and about his health with an openness and honesty I never expected.

I knew all of this was special at the time, that I was unusually lucky to have such a good supervisor. I would often hear horror stories from PhD students in other departments and institutions about the difficulties they had with their supervisors, and I would feel embarrassed by my riches. But I never reflected on how the relationship itself had changed. It is only now that he has gone, that I can see that he was not just a supervisor but a mentor and a friend. I only wish I had recognised this sooner and that I had been able to thank him whilst he was alive for everything that he had done for me.

I have spoken to many people about Philip in the last few months, who he was, what he meant to me and the hole his loss has left. As amazing as the support that I have received is, I have been left wanting to share my thoughts, feelings, and grief with someone who could truly understand the peculiar nature of the relationship. I am only now realising that the person I want to talk to about Philip's death, who I know would understand and would have the right words, is Philip. I think that speaks volumes as to his character and the uniqueness of his support and role in my life.

One of the first things I ever remember Philip convincing me of, was how a life can improve even after one's death. In classic Philip style, he spoke of a racing car driver (Philip would be dismayed that I can't remember who) who died during the end of the F1 season. Despite the driver's death he won the championship on the last day of the season as those chasing his point tally could not catch him. Philip argued that the driver's life had gone better due to posthumously achieving something that he had invested so much in. Philip was so heavily invested in his students, and his impact so strong on us, that I take comfort in knowing that Philip's life will continue to improve for some time to come.

With that in mind, I dedicate this thesis to Philip, the completion of which is as much his achievement as mine.

Contents

Acknowledgements	3
Dedication	5
Contents	7
Introduction.....	9
Motivation.....	10
Methodology.....	20
Assumptions.....	24
Aims	32
Overview	33
Chapter 1 – Horton’s Dilemma.....	37
Introduction	37
First Horn	39
Second Horn.....	44
Summary	57
Chapter 2 – Local Relevance Views and their Matching Principles.....	60
Introduction	60
Strongest Decides.....	63
Sequential Claims Matching.....	67
Hybrid Balance Relevant Claims.....	71
Match in Order of Strength.....	74
Match in Reverse Order of Strength.....	77
Concluding Remarks	78
Chapter 3 – Match by Closeness	79
Introduction	79
Matching by Closeness	80
Independent Justification	85
Testing Match by Closeness.....	90
Ambiguity.....	96
Alternatives	106
Concluding Remarks	110
Chapter 4 – Tiebreaks and Two Types of Relevance	112
Introduction	112
A Trilemma.....	117

Two Types of Relevance.....	120
Amending Sequential Claims Matching by Closeness	123
Deeper Justification	129
Conclusion.....	133
Chapter 5 – Two Approaches to Tiebreaking	136
Introduction	136
The Relevance Approach	137
The Outweighing Approach	145
Chapter 6 – Aggregate Relevant Tiebreakers	167
Introduction	167
A Simple Solution	168
Deeper Justification	172
A Complete Account of Aggregation	177
Conclusion	193
Appendices	197
Appendix 1 – Risk, Reductio and Weak Additivity.....	197
Appendix 2 – Sequential Claims-Matching by Closeness – Shortened	207
Appendix 3 – Fairness Last.....	208
Appendix 4 – Match in Order of Strength.....	211
Appendix 5 – Competing Tiebreakers.....	219
Appendix 6 – Local vs Global Relevance for F-claims	221
Appendix 7 – When are F-claims irrelevant?	226
Appendix 8 – Fair procedure	230
Appendix 9 – Tiebreaking and the Principle of Net Addition.....	238
Appendix 10 – Aggregate Relevant Tiebreakers and Expected Value	241
Appendix 11 – Partial Leftovers	243
Bibliography.....	248

Introduction

The world has finite resources, we have infinite desires. This is the fundamental economic problem as described to me on my first day as a fresh-faced undergraduate. Essentially, there are only three responses to this problem. First, we can maximise our resources. This is the response of the scientist, engineer, and industrialist. Second, we can limit our desires. This is the response of the priest, monk, and spiritual leader. Third, we can prioritise our desires and allocate resources accordingly. This is the response of the economist, politician, and ethicist.

This thesis is the work of an ethicist and so we will be exploring this third path. Crucially, this does not constitute a critique of the other approaches. The ethicist cannot solve the fundamental problem, at least not by herself. Only the first two approaches can truly solve the problem and align desire and resource: either we extract unlimited resources, or we limit our desire. Despite the immense progress made by the scientist, engineer, and industrialist - and living conditions that would make a medieval king jealous - we still seem no closer to satiating our desires. Nor has the work of the priest, monk, and spiritual leader ever been able to fully eliminate desire. Of course, we ought to read the disjunct here as inclusive. The only solution must surely be to limit (but not eliminate) our desires and extract just enough resources to fulfil them.

Nevertheless, whilst we wait for the industrialists and mystics to find middle ground, the ethicist will work. Her work is only valuable as far as we have no solution to the problem, being fundamentally focused on the non-ideal. How then ought we to distribute our limited resources? Unlike the economist and politician who primarily focus on the practical reality of distributing resources, the ethicist has freer rein to operate without such complicating factors. Instead, she focuses her efforts on deciding the right way to distribute said resources. Her contribution is thus to use ideal theory to shine a light on the non-ideal, in the hope of bringing the non-ideal closer into line with the ideal.

Motivation

This ethicist wants to focus on one particular sub-question of how to distribute scarce resources.

Namely, how do we distribute resources when those resources could provide large benefits to a few or small benefits to many: ought we to benefit those with the most at stake or ought we to maximise the benefit no matter how it is distributed?

This thesis will develop and defend a version of a view, called limited aggregation, which argues we ought to sometimes benefit those with the most at stake and sometimes ought to maximise the total benefit. The thesis tries to capture and explain our common-sense morality and use our judgements about simple cases to inform our approach to much more complex cases. Along the way we will consider issues of priority, partiality, fairness, and respect for persons.

Motivation

The first, and perhaps most important, task of the entire thesis is to explain and motivate the problem. Before we can look at whether my answers are any good, we need to know whether the question is worth asking. To do so I am going to first look at why this issue is philosophically interesting and then I will address why it is practically important.

The Problem

Firstly then, what is aggregation? What is the problem under consideration?

The easiest way to understand it is to consider an example:

Charity: A distant aunt has recently died and left her entire estate to charity. However, she did not specify in her will which charity it ought to go to. As her only living relative you are tasked with deciding how to spend the money. You narrow down your options to two choices. You could give the money to a charity that would use the money to perform a large number of cataract surgeries for people in the developing world. Or you could spend the

Introduction

money on a dialysis machine for a rural hospital in a poor region, which would allow them to extend the lives of a small number of kidney failure patients.

What ought you do? Do you save the lives of the kidney failure patients, even though there are less of them? Or do you improve the vision of many more, even though the benefit to each person is much less?

To decide we need to know how to compare the benefits of the two groups. This typically involves considering the benefits to each individual who would be helped and adding up all the benefits of both options. The option which generates the most total benefit on aggregate is the one we ought to do. This approach is broadly quite simple (once we agree on the benefit of a particular intervention) and can be very useful in deciding what we ought to do.

However, problems arise when the amount of benefit and the number of beneficiaries between the options differ greatly. Consider the following case:

Cough Sweet: When distributing your aunt's resources, you narrow down your options to two choices again. You can buy the dialysis machine, or you can buy a huge number of cough sweets and distribute them for free to anyone who develops a sore throat.

In this case, it seems that the aggregative approach might lead us to the wrong decision. If there are enough people who would benefit from the cough sweets, then this option could outweigh the dialysis machine. But this seems wrong to many - it seems that no matter how many cough sweets you could purchase, you ought to buy the dialysis machine instead.

This is the issue at the core of the aggregation literature. Is our intuition right in *Cough Sweet* and if so, what explains it? If we are to explain this intuition, we need to understand when aggregation is permitted and when it is not. This is philosophy though, so of course there are a number of further complicating factors that turn this seemingly simple question into a moral minefield.

Motivation

Broadly speaking there are three aggregative positions. Pure-aggregation, anti-aggregation, and limited aggregation. Pure-aggregationists say that we always ought to aggregate - in other words they maintain that the quantity of benefits always matters. Anti-aggregationists tell us never to aggregate - they might take a number of different approaches, the most common is to say that we ought to treat whoever has the strongest individual need (or claim) - in other words they maintain that the quantity of benefits never matter, and we should only be focused on the quality.

Limited aggregationists (or partial aggregationists) try to take a middle route - they say that sometimes aggregation is appropriate and sometimes it is not - in other words they maintain that quantity sometimes matters and sometimes does not. This approach is, for many, the most intuitive; on a practical level, it seems to get the right verdict in the greatest number of cases, and on a theoretical level it seems the most modest when it tells us that sometimes quantity matters and that sometimes it does not. However, it is theoretically the most difficult to understand, explain and implement - especially when our intuitions are complex and seem to conflict.

It is this limited aggregation position that I want to develop and defend in this thesis. I hope to demonstrate that some of the strongest arguments against the theory fail; that seemingly conflicting intuitions are actually consistent and can be captured if we make a number of important distinctions; and I hope to uncover several other intuitions and extend the theory to capture these intuitions.

Philosophical Interest

Now, let us look at why this is a philosophically interesting question. The easiest way to do that would probably be to point out how much recent work has been done on the topic and the number of impressive philosophers who are addressing the issue. However, let us not start the thesis with an appeal to authority, and instead actually look at why it has garnered such interest.

Introduction

Aggregation is an interesting topic in and of itself for a number of reasons. Firstly, it is an interesting battleground for discussion between consequentialist and non-consequentialist approaches to ethics. For consequentialists, aggregation is a simple matter. We ought to do the action which brings about the most good, no matter how it is distributed. Thus, we simply add all the various benefits and burdens on each side and do the thing that has the most total weight of reason - even if this is distributing cough sweets rather than funding dialysis machines.

Non-consequentialist accounts, which are often motivated by our intuitions around these aggregative cases, obviously want to resist such simplistic accounts of what we ought to do in these situations. However, settling on a positive answer for what ought to be done is more complex than just arguing that the consequentialists are wrong, and there is much disagreement in this regard. Of course, if non-consequentialists struggle to find a good account of aggregation, or all the options are shown to have terrible consequences, then this weakens the non-consequentialist position and thus strengthens the consequentialist one. Thus, the aggregation debate is an important battleground between these two central ethical theories.

It is also philosophically interesting because of how complex our intuitions are in the cases. We have many seemingly conflicting intuitions, and subtle changes in cases can alter drastically what we intuit we ought to do. This complexity indicates a normatively rich environment, with lots of interacting features. Those who take intuitions seriously will want to find a way to make sense of them in these cases and explain why they seem to conflict and what subtle features we are picking up on. Thus, there is much to explain and great potential to identify previously unnoticed moral features or interactions. On the other hand, if such progress cannot be made, and our (seemingly) conflicting intuitions are just inconsistent, then these cases will function as a mark against intuitionist approaches.

Practical Importance

Of course, part of what makes the topic philosophically interesting is also what makes the issue so practically important. In this section I want to quickly cover some of these issues. The most obvious issue is the one that the thesis deals with most directly. Namely, how we ought to distribute limited resources. Of course, setting health care priorities is the most salient example of cases when we must decide between different groups to save but it is not the only one. Other cases exist in which we might want to resist the consequentialist approach and say that certain considerations should not be taken into account in light of other, much stronger considerations. Cases which are getting increasing attention, especially among effective altruists, are issues regarding sentience, longtermism and, perhaps more idiosyncratically, decision making AI.

Healthcare

Firstly, let us look at the most obvious issue and the one the thesis most clearly grapples with. In recent years there has been an increasing focus on the fact that there are limited healthcare resources and that these cannot be distributed on a first-come-first-served basis. Difficult choices have to be made and these cases obviously involve aggregation. These serve as paradigmatic cases of aggregation and will form the basis of the examples I use throughout the thesis.

On a very basic level, sometimes there are choices that healthcare institutions and practitioners need to make with their resources. For instance, an institution might need to decide between funding surgery for multiple people with painful bunions, or funding more expensive cancer treatments for fewer people. Or a surgeon might need to decide between attempting a complicated lifesaving procedure on one patient or attending to two patients who could be saved from paraplegia in the same time. How we weigh the reasons in these cases is vital in determining what we ought to do, and issues of aggregation are obviously vital to making these decisions. My thesis, in shining a light on aggregation, might then improve how we distribute such resources.

Introduction

To expand on this a little, a more specific example of the importance of the aggregation literature for setting healthcare priorities is the Global Burden of Disease (GBD) tool (Murray & Lopez, 1996). GBD is a tool for measuring the total health impact of all diseases, injuries and other health issues.

Understanding the health burden of each disease, how diseases impact health, and the biggest health issues in each country makes GBD an invaluable tool for policy makers in deciding how to set their healthcare priorities. Aggregation is of particular importance to this literature, because the big picture considerations that limited aggregationists are considering are present here. The complexity of the issues at this level goes significantly beyond the norm, and policy-makers have the time and resources to properly consider all the features in a way frontline doctors and nurses do not. This means that complex aggregation considerations are particularly relevant and unusually helpful.

Now there are two ways in which the aggregation literature can impact GBD. Firstly, aggregation literature might impact how we measure the GBD, and secondly the literature might change how we use the GBD to make healthcare decisions. As GBD advocates are happy to admit, there is no purely descriptive way to measure GBD, and normative considerations are naturally baked into the tool (Murray & Lopez, 1996, pp. 7-13). Thus, limited aggregation considerations, and by extension my work, might have some impact on how we measure GBD. Issues of relevance and discounting might alter the way in which we calculate the GBD. However, I think there is limited utility to applying limited aggregation ideas to this issue. This is because doing so requires taking an axiological approach to limited aggregation. As I outline below, this is not uncontroversial. It would involve making the axiology and deontology coincide and would have the impact of making the GBD even less descriptive. Doing so might also throw up some odd conclusions where we improve the lives of some without making anybody worse off, yet in doing so we increase the GBD as previously discounted illnesses are now counted. Thus, it strikes me as best to avoid this approach to GBD.

Much more important, however, is how GBD is used to make decisions. On a pure-aggregationist approach, GBD measurements might be used to determine which treatments reduce the total

Motivation

burden of disease the most and funding such treatments. This is how GBD is currently used in practice. However, such an approach, of course, does not take into account the various considerations identified by limited aggregationists. If giving paracetamol to cure headaches reduces the GBD by more than funding expensive treatments for MS patients, then the simplistic model would suggest curing the headaches. If we think this is wrong, then there is obviously a high impact question of how we should use GBD data to set our healthcare priorities.

Limited aggregationist and relevance approaches to setting healthcare priorities are already being considered by some institutions and are under regular review. For instance, the Dutch have considered limiting the allocation of their healthcare resources to only substantial burdens and ruling out minor burdens (Voorhoeve, 2020). Similarly, the state of Oregon has a prioritisation system for allocating its Medicaid budget. This system works as a prioritised list with the Oregon Legislature deciding which conditions are worth treating, and which treatments do not meet the threshold (Perry & Hotze, 2011). Even in the UK, famous for its resistance to explicit healthcare resource prioritisation, there have been recent curbs on the prescription of over-the-counter medicines (NHS England, 2017) and the establishment of bodies to explicitly consider the funding priorities in local areas. These are quite rough and ready systems and there is quite a large scope for improving the theoretical tools and resources that these bodies have at their disposal. I have already begun to explore these implications with the Thames Valley Priorities Committee, and will be using insights from this thesis in the development of a new 'Ethical Framework' for the reorganised committee (Ethox Centre, 2023).

Importantly, I think there are two ways that my thesis can help address these issues surrounding healthcare and limited aggregation. Firstly, the thesis defends the limited aggregative approach against a number of objections and shows how limited aggregation can consistently capture our intuitions in an overall coherent manner. Doing so strengthens limited aggregation approaches and shows that they are good candidates for the right answer and should not be dismissed from

Introduction

consideration. This element of the thesis might only affect the credence in what we believe rather than the detail of what we believe, but this might still be sufficient for it to have significant impact. If the arguments I provide can persuade people of the virtues of such approaches, then those approaches are more likely to be used in the cases outlined above. This is an important consequence. For instance, in Chapter 1 I defend relevance views against a seemingly fatal dilemma which might be thought to render limited aggregation views unusable. By solving the dilemma, I remove the worry that might prevent the use of limited aggregation views in real world decision-making scenarios.

The other way that my thesis might help is in the detail. The approach to limited aggregation I take differs from the approaches taken by others: my approach adds detail and uncovers principles previously undiscovered and thus handles some of the cases differently. Of course, this changes the way we ought to apply our ideas in real life cases too, improving the way in which we distribute resources and giving us more analytical tools to approach difficult cases with. For instance, in Chapter 3 I defend a matching principle I call *Match by Closeness*. This principle tells us the order in which we should consider claims when we have to make difficult choices and thus helps us to decide what we ought to do in practical cases. In this way, my work improves the limited aggregation thesis as well as defends it.

Wider Implications

Of course, setting health care priorities is the most salient example of when we must decide between different groups to save, but it is not the only example. Other cases exist in which we might want to resist the consequentialist approach and say that certain goods or claimants should not be taken into consideration in light of other much stronger considerations. Two such cases, which are getting increasing attention (especially among effective altruists), are issues regarding sentience¹

¹ See Singer (1990) as the founding text for treating non-human animal goods equally to human goods, and Kagan (2019) for a 'hierarchical' account of animal ethics.

Motivation

and longtermism². In both cases, if we want to avoid some extreme consequences when aggregating many smaller goods, we might want to apply some sort of discounting rules or principles around when goods are relevant.

In the longtermist case, whilst we might want to say that future lives are worthwhile and give us certain reasons to act, we might think that there is a limit to what such considerations can outweigh. For instance, we might want to say that a life in a thousand years' time is simply not relevant to a life today. Perhaps we might want to say that the life in a thousand years' time is as valuable as a life today, but that the reason provided by the current life is so much stronger given its proximity that it disables the reason provided by the life in the future.³

Similarly, in the sentience cases, we might want to say that lives much further down the sentience scale might not be relevant to lives much further up the scale. Or at least the harms or moral claims of those further down the sentience scale might not be relevant to the harms or claims of those further up the scale. This admits of a certain amount of flexibility. We might say that for life saving treatments the lives of creatures significantly further down the scale of sentience might not be relevant and thus their use in medical science might be permissible. Conversely, the small amount of pleasure derived from eating meat is unlikely be a relevant factor compared to the lives of such animals.⁴

Lastly, and perhaps more idiosyncratically, the thesis has some potential implications for decision making AI. As I am at pains to point out in the below section, the main focus of this thesis is to describe a theory of the right, rather than to provide a decision procedure. The sequence that I

² For a good overview of longtermism see Greaves & MacAskill (2021). For the beginnings of a response, that might be compatible with, and could be strengthened by, limited aggregation, see Mogensen (2019).

³ This will be further complicated by the introduction of the non-identity problem, whereby the person whose life you could save in the future is not the same as the person who would lose their life early if you do not act. Taking into consideration the partial concern or personal perspective of the individuals in such cases would be much harder, if not impossible. How this relates to relevance views is an important and interesting question that I will not be addressing in this thesis. Of course, how it relates will depend on how we understand relevance views in simpler cases and so making progress on relevance clears the ground for this work.

⁴ For a recent account that explores limited aggregation in human-animal relations see Egel & Martin (2022).

Introduction

describe in the later part of the thesis is quite complex and thus would not be easy to use as a decision procedure in real life cases, especially with the number of other confounding factors. Nevertheless, some of the principles I do identify, such as Match by Closeness, are usable principles and there will be a number of rules of thumb derived from the sequence that will be usable. However, Artificial Intelligence may be a special exception in this regard.⁵ The sequential and algorithmic approach I take is amenable to use by AI and might help with AI decision making. Unlike human decision makers, AI may also have sufficient computing power to implement such algorithms quickly and easily.

However, the way in which the thesis might be relevant for AI is a bit more subtle than simply applying the algorithm I uncover here. Where there might be some success in doing so, the scope is limited and the number of moral features I consider in these cases is also limited. So naturally, other features would also need to be accounted for. Of course, the AI would also have to be supplemented with human decisions about relevance and how different considerations balance each other out.

The real implications for AI then, are in recognising that the fundamental structure of reasons (the theory of the right) and the way in which we ought to weigh reasons and decide what we ought to do might be best explained via algorithm. This would mean that such decisions are fundamentally amenable to AI. The sequence might then provide a prototype for a wider project of developing non-consequentialist/non-maximising decision procedures for AI, in contrast to current AI systems. I do not want to overstate my case, what I have here is all very tentative of course, and I shall not be delving too deep into these connections. Nevertheless, I believe there is value in mentioning these connections.

⁵ The other exception might be large institutions who have the time and resources to consider how best to distribute their resources, and already do so. I am thinking here about government bodies such as NICE in the UK and charity evaluators such as GiveWell.

Methodology

Next, I should address the methodology I will be using to consider the issue. Essentially, I shall be adopting Frances Kamm's methodology as she outlines in the introduction to Volumes I and II of her "Morality, Mortality" often described, as per the title of another of her books, as intricate ethics. In brief, Kamm's position is this. Rather than start with theory and see what consequences the theory has in particular cases, we start with particular cases and allow this to inform our theory. In complicated cases this means working through our intuitions in stages, rather than try to reach an immediate intuitive judgement of the entire case.

But before we get into exactly what this entails, let us first look at the variety of methodological approaches to see exactly why I choose the one I do. Kamm describes three distinct methodologies, we might call them the Theory-First approach, the Reflective Equilibrium approach and the Case-First approach (Kamm, 1998, pp. 5-9). Whilst Kamm does not articulate it in this way, it strikes me that we might understand these different approaches not as wholly distinct methodologies, but as different points along a spectrum.

On one end of the spectrum are the Theory-First style approaches. These approaches start with a well-defined ethical theory, usually come to through argumentation or rational reflection on core principles, and then they apply the theory to particular cases. Extreme versions of this approach will involve no consideration of particular cases at all; the theory is simply to be applied in those cases and the resulting answers to be accepted, however counter-intuitive they might seem.

On the other end of the spectrum, we have the Case-First approaches. These approaches look primarily at particular ethical cases, by considering our pre-theoretical judgements or intuitions in these cases they attempt to uncover our moral reasons. Only by describing and explaining our moral reasoning can we come to a better understanding of the various reasons in play and how they interact. Extreme versions of the Case-First approach would completely avoid theory, starting and ending with our intuitions or pre-theoretical judgements in individual cases.

Introduction

In the middle of the spectrum, we have the Reflective Equilibrium style approaches.⁶ Typically, we think of these approaches as starting with a theory and then amending the theory in light of particular cases. Where pre-theoretical judgements conflict with the theory we know we need to adjust the theory and there may be some back and forth between the theory and the judgements. Procedurally this is how these style approaches tend to work, but the true centre of the spectrum would give the theory and our pre-theoretical judgements equal weight. The theory and judgements would then be treated as different data points that need to be brought together. We might understand this approach as a coherentist project.

There are, of course, advantages and disadvantages to every position on the scale. Theory-First approaches are less reliant on intuition, though also potentially more susceptible to criticism based on counter-intuitive cases. In this way such approaches are quite rigid. This is both an advantage and a disadvantage. Positively, it means that the view has a strong theoretical basis and so if one is convinced by the theoretical underpinnings, then the aggregation view that falls out of it will be quite resistant to criticism - the proponent has a strengthened ability to bite bullets by pointing to the deeper theoretical justification. Negatively, however, it means that the view has less flexibility to cope with criticism and to be amended in light of counter-examples or difficult cases. If the counter-examples or difficult cases are too difficult these bullets might become very hard to bite and the theory as a whole might be threatened. In this way these approaches are well suited to handling slightly counter-intuitive implications but less suited to handling much more problematic cases.

Conversely the further along the spectrum we go the better suited the approaches are to coping with more problematic cases. Such approaches can make significant amendments to the theory in light of problem cases as and when necessary, or use problem cases as starting points from which we might learn more about the underlying structure of morality. However, these approaches struggle more with smaller bullets or cases where our intuitions might conflict or be hard to identify.

⁶ For a more detailed overview of Reflective Equilibrium see (Daniels, 2020).

Methodology

Even slightly counter-intuitive examples might thus need explanation. In such cases, it will be harder for these approaches to decide which of the bullets to bite when intuitions conflict, or to decide what to do when intuitions are hard to identify. Without a fixed theoretical justification, it will be hard to come down on one side or the other.

Furthermore, because such approaches will need to amend their views in light of smaller issues, the view will become more complicated. More importantly, it also means there is more opportunity for error to creep in and a greater chance that we will come to discover competing intuitions that might threaten the entire theory. If there are occasions when the view cannot be successfully amended in light of counter-examples, then such approaches are in trouble.

Now, I ought to indicate where on this spectrum my methodology lies. On a scale where the extreme Theory-First approaches are 0, the extreme Case-First style approaches are 10, and a perfectly balanced Reflective Equilibrium approach is a 5, my approach would be about a 7 or 8. I, like Kamm, then prefer the Case-First style approaches.

I think we should trust our pre-theoretical judgements or intuitions unless we have good reason to think we ought not to. In the absence of defeaters then our intuitions are trustworthy. Yet I do not want to say there is no theory underlying the structure of our reasons and that we need only rely on our judgements in each particular case.

I diverge from Kamm somewhat when it comes to our confidence in our intuitions in complex cases. Kamm is happy to “rely on intuitions even at great levels of complexity” (Kamm, 2007, p. 5) whereas I am a little less trusting of intuition in such cases.⁷ It strikes me that even with a great deal of concentration there are cases where it is impossible for us to see the whole complex landscape and

⁷ I find her argument a bit weird in this regard. She says we ought to start with artificial cases as we can better identify moral features or potential principles from these cases, but then goes on to argue that the consideration of such principles is not much help in the real-world cases, and that we are better off relying on our intuition even in such complex cases (Kamm, 1998, p. 9). If this is the case, then identifying moral principles becomes a purely intellectual pursuit with little practical upshot, which does not do justice to Kamm’s work.

Introduction

keep in mind all the various moral features and the detail in how they relate. If our intuition is to be understood as a moral sense then, unless this faculty is completely unique amongst our senses, there must be a limit to our abilities in this regard.

This also seems borne out by introspection. It seems we have intuitions about which group we ought to save when the groups are small, for instance, when we must choose to save either one person from dying or twenty from quadriplegia. But when the groups are much larger it becomes impossible to grasp all the factors needed to be able to generate an accurate pre-theoretical judgement. No one has a strong intuition about whether we should save Group A where we can save twelve from death, eight from quadriplegia, forty from lost fingers, ten from blindness and a thousand from headaches or Group B where we can save eight from death, twenty from paraplegia, ten from lost limbs, sixteen from deafness and two-thousand from sore throats.

Yet, I do not think this is a problem for Case-First approaches. In fact, one of the main advantages of this style of approach, that Kamm seems to underplay, is how it handles these cases. Using intuitions to identify principles in more basic and artificial cases allows us to have a basis to work on in much more complicated cases - cases where there is too much going on for us to realistically have any strong intuitions. By working through our intuitions in the basic and more artificial cases we can identify the various important moral factors and reasons as well as the structure of these reasons and the principles that underline them. We can then apply what we have learnt about the theory to the much larger and more complicated cases such that we do not have to rely, at least so strongly, on our intuitions in those cases.

One consequence of this style of approach is that we are likely to end up with a very complex theory. This is often seen as a downside. Analytic philosophy since, at least, Quine has become enraptured with desert landscapes and simplicity is considered virtuous. But, without wanting to be pulled into a discussion on the state of academic philosophy, I think this is a mistake.

Assumptions

The presumption of simplicity in the normative realm strikes me as profoundly odd. Normativity would stand alone as the only non-complex part of our world. Everything about the world, all of the sciences, arts and human interactions, point to a world full of complexity with the interplay of many different features. Our default assumption then should be that ethics is likely to be deeply complicated in the same way as every other facet of human life. When we look at any real-world ethical decision, we can see that this is borne out in reality too - many ethical problems we have to deal with seem so complicated as to be almost impossible to get right. In many ways, this *is* rocket-science!

Lastly, there are a couple of other advantages to this style approach. As Kamm indicates, this approach “permits recognition of new factors that may be morally relevant in certain cases, factors emphasized by no theory yet developed” (Kamm, 1998, p. 6). If we start with theory, then we are likely to explain the moral phenomena in such a way that they fit the theory, even if this means squeezing square pegs into round holes or running roughshod over complicating details. However, if we start with the cases, and our pretheoretical judgements, then we are less likely to miss these important details. Also, using artificial cases allows us to hold steady all the other factors (like scientists testing under controlled conditions) and investigate the difference a factor makes in isolation. This means we can be sure of the effect of the moral feature under consideration when we might otherwise ascribe that effect to some other feature.

Assumptions

I should now make it clear what I will and will not be addressing in this thesis. Firstly, I want to make very clear the scope and focus of this project. As my methodology indicates, the approach I am taking is fairly narrow, I am not envisaging myself to give a unifying theory working from first principles and building a theory of limited aggregation on top of these principles. Thus, there are many theoretical debates that I want to stay away from.

Introduction

Firstly, a few smaller items to note. I should make very clear, in case it was not obvious, that all the cases I discuss in the thesis are simplified examples. Thus, when I speak of comparing death against lost fingers, we should hold everything else equal such that we are not comparing the death of a centenarian with the lost finger of a concert pianist. These features are obvious intensifiers and attenuators and would thus affect the relevance of the lost fingers, but to keep things simple we shall avoid introducing such features.

Importantly, my arguments, and the account I develop, are not reliant on the precise examples I use throughout the thesis. There is wide scope for disagreement about exactly what is relevant to what, but that discussion is broader than the thesis. However, it is helpful to use precise examples to illustrate the points I make. I have tried to make these relatively uncontroversial, but inevitably not everyone will agree on whether X is relevant to Y. Thus, wherever necessary, the reader might need to substitute the examples with ones which better match their commitments. So long as the reader can find acceptable replacements, my arguments should not be affected.

I should also mention that some of the cases I use in the thesis have names, whilst others are only given numbers. For the sake of clarity across chapters, important cases, i.e., cases that demonstrate a central intuition or important problem, will be given names. To keep things simple and each of these important case names memorable, all other cases will simply be numbered.

A distinction it is also important to note is between interpersonal and intrapersonal aggregation. Interpersonal aggregation is aggregation between the claims of different people. Intrapersonal aggregation is aggregation of factors within a single person's life. Importantly, what might be permissible intrapersonally might not be permitted interpersonally. For instance, it might be permissible (although probably unwise) for someone to take a tablet that cures their sore throat instead of one that saves them from a one in a thousand chance of dying. But it is not permissible, to give a tablet to someone to cure their sore throat when it could be given to someone else to save them from a one in a thousand chance of dying. Intrapersonal aggregation is an interesting topic,

Assumptions

and any full account of aggregation will need to address it, as well as how intrapersonal and interpersonal aggregation relate. In this thesis however, I am only concerned with interpersonal aggregation. I simply do not have the space to cover both issues in detail.

Claims and Claimants

I should mention that throughout the thesis I shall talk in terms of claimants and claims. Now I primarily use the terminology of claims for the sake of simplicity. The term is common in the literature and is a helpful way to capture the 'location' of the moral reasons. In this literature, if somebody can benefit from a resource, then that person has a claim to that resource and we have a duty to give her that benefit. (Voorhoeve, 2014, p. 66) Similarly, if someone can avoid a harm by my help, then they have a claim to my help. In this way, the moral reason is explicitly tied to that person. The greater the benefit conferred, or harm avoided, the stronger their claim is. Thus, someone who will die without my assistance has a stronger claim to my assistance than someone who will simply lose a finger.

For the sake of the thesis, I shall be setting aside non-welfare considerations, such as desert or prior-commitments and relationships, which might affect a claim's strength. For instance, we tend to think that a son has a stronger claim to help from his father than the neighbour's child does because of their important relationship. Of course, these are morally important features and can provide their own anti-aggregative reasons. Any complete account of what we ought to do with our resources, both as individuals and as groups, will also need to include these considerations. However, the aggregation question that I want to focus on is not how the strength or relevance of a claim might vary according to these features. Instead, I want to focus on the more fundamental issue of how the strength or relevance of a claim might vary according to its relationship with other competing claims. Thus, we must hold constant these other variables. Focusing only on the welfare aspects of a claim will help us to do this.

Introduction

Furthermore, whilst I am partial to claims talk, I understand that some philosophers are not as keen as they “prefer to avoid the language of claims, as that language is most comfortably used in rights-based or contractualist theories” (Tadros, 2019, p. 175n2). Thus, it should be noted that I am not committed here to a rights-based view where any claimant has a particular right to aid or assistance. Moreover, my argument does not rely on claims talk. Thus, if one objects to claims talk, claims can be substituted for ‘duty-grounding facts’, and the relative strength of claims can be described as the difference between the duty-grounding forces of such facts (Tadros, 2019, pp. 172-178).

Lastly, to keep sentences shorter and easier to read, I shall talk of ‘death claims’ and ‘lost hand claims’ etc. A death claim is simply a claim against dying, or a claim to a resource that will prevent death; if I am dying and you have the cure, then I have a death claim.

Theoretical Explanations of Relevance

One debate that I want to stay away from in this thesis is about how exactly the notion of relevance works and under what precise conditions claims become irrelevant. There are a number of different theoretical explanations of relevance, which pick up on different features of these cases and emphasise varying moral features:

Alex Voorhoeve parses relevance out primarily from the personal perspective of the weaker claim: if the stronger claim is so much stronger that the weaker claimant could not press her claim, then it would be disrespectful to consider the weaker claim. If, on the other hand, the weaker claimant could licitly choose to save herself over the stronger claim then the weaker claim is relevant, and it would be disrespectful to ignore it (Voorhoeve, 2014, pp. 70-75).

Frances Kamm rejects understanding relevance from the personal perspective of the weaker claim (Kamm, 1998, pp. 170-171) (Kamm, 2015), and offers instead a number of different ways to incorporate the personal perspective (Kamm, 1998, Chapters 8-10). Kamm’s approaches put a

Assumptions

greater emphasis on the personal perspective of the stronger claim, but also on what is salient to a disinterested third party.

Victor Tadros denies that partial concern provides the right explanation for these respect conditions at all. Instead, he argues that such respect conditions are ‘morally basic’ and constitute “a plausible instance of the more general phenomenon of contextual interaction” (Tadros, 2019, pp. 177-178).

It might seem surprising (even irresponsible) for a thesis on limited aggregation not to address this debate, but I have good reason to stay away from it. My reasons are several-fold. Firstly, the whole point of the methodology I use is not to commit to any principles until one has considered all (or at least a large number of) the particular cases. The view I describe and the cases I look at might rule out some of these theoretical explanations, so it would be premature to attach myself to one account. Secondly, and more importantly, the view I expunge is amenable to different explanations of how claims become irrelevant, and so I need not commit myself to any particular justification. As such I will rely on common sense conceptions of irrelevance and wherever possible will make my examples ones that would be uncontroversial to all limited aggregationists.

However, since partial concern and respecting the personal perspective of each of the claimants is a common thread through most of these views, then whenever it is helpful to further develop the deeper justification of my view I will talk of these aspects, though not necessarily with any fleshed-out detail. This will also help to show that the distinctions I make and the view I express are not ad hoc. When I think it is particularly important, I will also indicate how the view developed can be captured or understood in terms of particular theoretical underpinnings. I hope this will help to ground the views I expunge and also make clear how the developments I make should be acceptable to those aggregationists whilst avoiding committing to any particular approach.

There are a number of reasons why I think this is the best approach. Firstly, it is the structure of the reasons here that I care about - not the specifics. Focusing on the specific justifications will bring in unnecessary controversies and distract from this core aim. Secondly, I am not trying to convince

Introduction

those who disagree, but merely explain our intuitions and show how they are all consistent.

Committing to a particular theory of relevance is unlikely to help in this regard. Thirdly, I do not want to narrow down who might accept my ideas. By developing my view in the way that I do, should make my conclusions acceptable to as many limited aggregationists as possible.

Fourthly, I am trying to work a bit further downstream and to create something potentially a bit more practicable. Whilst I do not want to say that the debate about the way in which a claim has relevance has no practical benefit *ipso facto*, I do believe it is less important than recognising the structure of how relevance works. Lastly, it strikes me that the literature regarding relevance is largely driven by our intuitions anyway. Thus, it is not a bad idea to stick to the intuitions, at least until we have them all consistent, before we start to theorise about the best explanation for a claim's relevance or irrelevance. If we start to do this latter aspect too soon, then we are likely to encounter issues further down the line.

Deontology and Axiology

Another question that I think is important to address early and set aside is the question as to whether the view I am developing is deontological or axiological in nature. I alluded to this distinction in my earlier discussion of limited aggregations potential impact on Global Burden of Disease. Again, there is a temptation to side-step this question as there might be disagreement as to whether limited aggregation is fundamentally a deontic or axiological thesis. For our purposes I think it is best to approach the issue in entirely deontic terms. This is because I think all limited aggregationists can broadly agree that the conclusion of any limited aggregation view is to tell us what we ought or ought not to do.

There is a separate issue as to whether the axiology matches the deontology, but I will set this question aside. Firstly, I do not want to narrow down who might accept my ideas. By treating the issue in deontic terms and leaving the question as to whether the deontology matches the axiology, what I say should be acceptable to all limited aggregationists. Those who prefer an axiological

Assumptions

approach would need only recognise that what I say in terms of what we ought to do could be matched by underlying values. Thus, when I say that a claim is irrelevant (say when a sore throat comes into competition with death) I am saying that the reason to save the sore throat ought not to be taken into consideration. This could easily be converted into an axiological thesis by also saying that the goodness of saving the sore throat is disabled, or defeated, in this context - that it would not be good to save the sore throat.

My other reasons for preferring this approach and avoiding this discussion are somewhat similar to those I mention regarding what makes a claim relevant or irrelevant. Again, I care more about the structures and interaction between the various phenomena and am less interested in making claims about exactly what the phenomena are, or how best to characterise the phenomena. I want to stay slightly further downstream than discussing whether the thesis is axiological or deontological, and I do not believe that seeing the project through primarily deontic or axiological eyes will make a difference to the practicable aspects of my work, nor will they effect the various principles and ideas I uncover.

Theory of Right and Decision Procedures

The next important distinction that I ought to cover is whether the project is essentially a project in the theory of right or a decision procedure. In brief, a theory of right answers the question 'what is the morally right thing to do?' whilst a decision procedure answers the question 'how do I decide what the morally right thing to do is?'

This is a much more complicated matter to deal with than the other two matters. Firstly, I think it is worth mentioning that the distinction between the theory of right and decision procedure is less clear on non-consequentialist and broadly deontological views. Deontological views have a theory of right very closely linked to *how* we ought, and ought not, to do things. Thus, we might think that the two do not come apart in quite the same way as they do on other views. Nevertheless, the purposes

Introduction

of a theory of right and a decision procedure do still come apart, and how my account will be judged will depend on what the project aims to be.

As will become clear in the following chapters, the view I develop involves creating a sequence, or procedure, to get to the right answer. It would be easy to assume that what I am creating is a decision procedure. However, this is not the primary goal of the thesis. The thesis should be primarily judged as an attempt to capture a theory of rightness and to demonstrate that there are ways of capturing intuitions that are seemingly contradictory.

The sequence I develop, whilst it looks like a decision procedure, is actually a complex non-additive theory of the right - it gets the right answer and shows how the various features interact. Thus, the best way to understand the sequences is as a complex algorithm that captures a holistic, but computable, approach to reasons.

Nevertheless, the view I develop might have some potential use as a decision procedure. I do not want to overstate my case here and will not be developing any arguments in regard to this, but there is still some scope to provide some practical decision theoretic tools. Firstly, some of the broader principles I aim to uncover will be practically usable. For instance, the *Match by Closeness* principle I develop in Chapter 3 can be used as a rule of thumb when comparing large groups even when the use of the entire algorithm is not appropriate.

Secondly, I hope that amended versions of the sequence itself might be usable in some limited decision-making situations. Situations where there are time and resources to adapt and implement the sequence and where its use might be particularly helpful. I am thinking here about large organisations such as the UK's medical authority NICE and EA charity evaluators such as GiveWell. I shall return to these potential upshots in the conclusion.

Aims

So, with our methodology and our assumptions in view then, let us state clearly the aims of the thesis. Most of these points have already been covered in some detail above, but it would be prudent to list them independently.

Firstly, I aim to uncover some new intuitions in aggregation cases. This might seem like quite a difficult task, as so much of the aggregation literature has already gone through most cases with a fine-tooth comb. Nevertheless, there are cases with subtle differences in which our intuitions seem to differ that have not been reported on in the literature. At the very least, these cases have not been considered in concert and the seemingly conflicting intuitions generated when we do, throw up serious questions for many of the current views. In this way, the first aim of the thesis (though not the first part) is to generate new data points from our intuitions, which need explaining, and to demonstrate that all current theories cannot capture these points. This, of course, then opens up space for new or amended theories that can explain this new data.

Following on from this, I aim to generate a new way to capture all these new data points, alongside the old data points. In keeping with a Kamm-style methodology, I want to find a way to explain all these intuitions, and only consider abandoning them if there is no reasonable explanation for them. In some of these cases, this might be quite a complex task as some of the intuitions might, on the face of it, seem contradictory or inconsistent. The aim then is to show that all these intuitions can be held consistently and that there are genuine reasons to distinguish between the cases. Showing that the intuitions are consistent will itself work as a defence of both the Case-First style approach and the veracity of our intuitions. Thus, it gives us intellectual permission to remain committed to our intuitions and the limited aggregation style approach that accompanies them.

The third aim is then to uncover the underlying principles that might explain and capture our pre-theoretic judgements. This aim is not completely separate from the previous aim of course, as part of showing how these intuitions are consistent with each other will involve giving theoretical

Introduction

explanations. Of course, any principles, or reasons, which are uncovered will need to be backed up with deeper justification, and this is something I also aim to provide. However, in keeping with the style of approach I outlined above, this will not involve connecting it to any particular theory. I shall do my best to explain the reasoning and the justification behind any principle in a way that does not commit me to any such theory. Doing so will, I hope, also show that the implications of my findings are wider than amendments to a single theory but might be applicable to a number of different approaches. Where appropriate, and where I think it will not be a complicating distraction, I aim to outline how the deeper justification to my view might be captured or understood through different approaches to limited aggregation.

Lastly, as discussed above, the focus in this thesis is on the theory of the right rather than on any decision procedure. Nevertheless, I hope to show some potential decision theoretic takeaways and will aim to develop the view in such a way that future research will be able to explore the decision theoretic potential of the approach. Similarly, I aim for my findings to be usable in applied ethics and to provide an entry point for more detailed further research into practical questions, specifically in population level bioethics.

Overview

Given the complexity of the arguments and the position I come to in each chapter of the thesis, I shall keep my overview of the thesis short. The following are quick summaries of the issues and approaches taken in each chapter. They should provide a guide to which issues are considered where and give an idea of the flow of the thesis. Section 1, consisting of Chapters 1 through 3, focuses on weighing claims of different strengths against each other and determining which group has the stronger overall claim to be saved. Section 2, consisting of Chapters 4 to 6, then looks at tiebreak cases (i.e., cases where the groups of claims are balanced), the requirements of fairness, and how we can amend our view to choose between balanced claims.

Section 1 – Competing Claims

Chapter 1 - Horton's Dilemma

In the first chapter I shall address the biggest problem for limited aggregation views. Before we move on to the details of my account, we must first ensure that there are no fatal issues for the view as a whole - that we are not building a house on sand. Joe Horton in his 'Always Aggregate' presents a 'fatal dilemma' for limited aggregationists. If there is no solution to the dilemma, then the rest of the thesis becomes redundant; I would be solving smaller disputes amongst wrong options, and there is not much use in that. Addressing the dilemma will inevitably inform how we approach limited aggregation in future chapters: doing so will uncover some of the underlying principles and assumptions to which I will (and will not) want my approach to conform.

In this chapter, I shall argue that the 'fatal dilemma' is neither fatal nor a dilemma. I shall show that the First Horn of the dilemma is devastating but avoidable and the Second Horn is unavoidable but not devastating. Nevertheless, I will show that Horton's argument does help us to narrow down the acceptable range of views.

Chapter 2 - Local Relevance Views and their Matching Principles

The focus of this chapter is to consider the strongest Relevance views so far offered and to show why they all fail to capture our moral thinking. Relevance views capture our intuitions in aggregation cases by matching or balancing competing claims, until one group has claims which can no longer be matched. Different Relevance views give different answers as to how (and in what order) we ought to match and balance the competing claims.

In this chapter, I consider the different ways we might match or balance competing claims. This is a question that has not been sufficiently addressed in the literature, yet Relevance views live or die by their matching principles. To do so we will survey all the strongest matching principles offered in the literature. I show that these matching principles all fail to escape Horton's dilemma, violate several commonly held principles and break the internal logic of Relevance views.

Chapter 3 - Match by Closeness

In response to Chapter 2, I develop a new matching principle which I call *Match by Closeness*. This principle prioritises matching similar claims first, and only compares dissimilar claims when other comparisons are unavailable. I show that this principle gives intuitive judgements in a number of cases. Furthermore, *Match by Closeness* avoids all the pitfalls of the other matching principles - it avoids the First Horn of Horton's dilemma and has a good explanation for landing on the Second Horn. Importantly, the principle is independently plausible and provides a deeper justification for the principles discovered by our case judgements. Match by Closeness also respects several other important principles and fits well with wider approaches towards limited aggregation by considering each claimant's personal perspective.

Section 2 – Tiebreaking Claims

Chapter 4 – Tiebreaks and Two Types of Relevance

In the fourth chapter, we will look at how Relevance views should treat tiebreak cases i.e., cases where weak claims might tiebreak between balanced stronger claims. I present a trilemma for two different broad approaches to Relevance: Local and Global Relevance. Namely, that neither approach can capture our intuition in tiebreak cases, without forfeiting our intuitions in other important cases.

Fortuitously, there is a way to salvage relevance views if we consider a Hybrid view that combines the two approaches. I demonstrate that by combining the strengths of both approaches we can capture our intuitions and depict the aggregation landscape more accurately. Plus, in keeping with the methodology previously outlined, I explore the deeper justification behind the Hybrid view, beyond its ability to make all our case judgements cohere.

Chapter 5 – Two Approaches to Tiebreakers

In Chapter 5 we address another important issue with tiebreaking cases, namely whether tiebreaking should be sensitive to the number of claims. According to Kamm, and most Relevance

Overview

views, tiebreaking should not be sensitive to numbers. Conversely, according to Broome's outweighing view tiebreaking should always be sensitive to numbers. In this chapter, I demonstrate that both views fail to capture some important intuitions. Namely, that when tiebreaking, sometimes numbers count and sometimes they do not. Plus, even in cases where both approaches get the right answer, they often do so for the wrong reasons.

Chapter 6 – Aggregating Tiebreakers

With the criticisms from Chapter 5 in view, I outline another new hybrid approach that combines the outweighing and relevance approaches into a view I call *Aggregate Relevant Tiebreakers*. *Aggregate Relevant Tiebreakers* captures all our intuitions in these cases and combines the justifications of both prior views in a way that more comprehensively captures our thinking about fairness. I will then bring all the work in the previous chapters together. I demonstrate how we can combine *Match by Closeness* and *Aggregate Relevant Tiebreakers* into one comprehensive account that covers everything we want to say about aggregation in cases of pairwise comparison. Importantly, this includes complex heterogenous cases with many claims of different types on both sides.

Part 1 – Competing Claims

Chapter 1 – Horton’s Dilemma

Introduction⁸

As we have seen in the introductory chapter, sometimes we must choose between competing claims to aid or assistance, and sometimes those competing claims differ in strength and quantity. In such cases, we must decide whether the claims on each opposing side can be aggregated. Limited aggregationists try to capture our intuitions in a variety of cases by sometimes requiring and sometimes disallowing aggregation. They argue that a set of claims can be aggregated only if they are sufficiently strong (compared to the claims with which they compete) to be morally relevant to the decision.

For instance, paraplegia is sufficiently close in strength to death that aggregation is permitted; it is not disrespectful to a dying man to treat a number of paraplegics instead of him and so the claims against paraplegia can be pressed. On the other hand, losing a hand is so distant in strength that it would be disrespectful to treat any number of people who will lose a hand instead of the dying man.⁹

Now, relevance views admit of two interpretations - Global Relevance and Local Relevance (Tadros, 2019). Both views handle homogenous cases, i.e., cases where each group only contains one type of

⁸ A version of this chapter is forthcoming in the *Australasian Journal of Philosophy*.

⁹ To reiterate, nothing hangs on these particular examples and the reader is encouraged to replace what is at stake here with other claims if she has a different standard for relevance.

Introduction

claim, in the same way. But the two views differ in heterogenous cases, i.e., cases where the strength of claims differ within groups as well as across groups.

Global Relevance views consider whether a claim is relevant by comparing that claim to the strongest claim with which it competes. If a claim is not relevant to the strongest claim with which it competes, then it can play no part in which group to save. In this way Global Relevance decides whether a claim is relevant or irrelevant simpliciter.

Local Relevance views, on the other hand, consider whether a claim is relevant to *each and every* type of claim with which it competes. So, a claim might be irrelevant regarding the strongest claim in the competition, but relevant to a weaker claim in the competition. That a claim is irrelevant to the strongest claim does not mean it cannot play any role in which Group to save. Thus, if the strongest claims are counterbalanced, then weaker claims can play a role in deciding which group to save. For instance, a lost hand claim might not be relevant to a death claim, but if that death claim is counterbalanced by another death claim, then the lost hand claim might still be able to play a role in deciding who we ought to save. Perhaps by tiebreaking between death claims, or by counterbalancing other (weaker) claims that might have otherwise been decisive.

Now, before we move on to the details of my account, and decide between Global and Local Relevance, we must first ensure that there are no fatal issues for limited aggregation approaches as a group. Otherwise, we are building a house on sand. Joe Horton in his *Always Aggregate* (2018) presents a 'fatal dilemma' for limited aggregationists in heterogenous cases. He argues that all limited aggregation views necessarily land on one of two horns. Views that land on the First Horn seem to suggest that a winning group that is strengthened by more claims of at least equal strength, could become the losing group, whilst the losing group strengthened by lesser claims can become the winning group. On the other hand, views that land on the second horn seem to suggest that combining two losing groups together and two winning groups together can turn the losing groups into the winning groups and the winning groups into the losing groups.

It is important to address this issue first for a number of reasons. Firstly, if there is no solution to the dilemma then the rest of the thesis becomes redundant - I would be solving smaller disputes amongst wrong options, and there is not much use in that. Secondly, how we address the dilemma will inevitably inform how we approach limited aggregation for the rest of the thesis. For instance, if we discover that only one of the horns is not fatal, then we know that this is the horn our view should land on. It would be an upsetting waste of time, if I were to develop an account that handles many of the other problems only to discover that there is an escape from the dilemma, but one that my approach cannot navigate. Furthermore, by tackling the biggest issue first we will be able to uncover some of the underlying principles and assumptions that I will and will not want my approach to conform to.

In this chapter, I shall re-examine Horton’s argument and demonstrate that the ‘fatal dilemma’ is neither fatal nor a dilemma. I shall show that the First Horn is devastating but avoidable and the Second Horn is unavoidable but not devastating. Nevertheless, I will show that Horton’s argument does help us to narrow down the acceptable range of limited aggregation views.

First Horn

Let us look at the Horns separately. Let us start with the First Horn of this dilemma. I shall use slightly different examples and a different style to those used by Horton, both for my own ease and to maintain a consistency of style throughout the chapter and thesis. The conclusions of his argument will not be affected. The First Horn is made up of two stages:

First Horn Case:

Stage 1 - You can save Group A or Group B. Group A contains a death claim. Group B contains one thousand lost hand claims.

Stage 2 - One hundred lost limb claims are added to Group A. Six lost limb claims are added to Group B

First Horn

Now for our purposes, let us assume that lost hand claims are irrelevant to death claims but are relevant to lost limb claims.¹⁰ Thus, no number of lost hands ought to be saved instead of a life, but some number of lost hands ought to be saved instead of a lost limb. Let us also assume that lost limb claims are relevant to death claims, and so a sufficient number of lost limb claims can outweigh the death claim. Let us also stipulate that four lost limb claims match exactly with a death claim, and thus five or more lost limb claims will outweigh a death claim.

For illustrative purposes I use tables throughout this chapter and thesis to help visualise these cases and make clearer exactly where the issues and solutions are.

First Horn Case

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia		
3 - Lost Limb	(100)	(6)
4 - Lost Hand		1000

For ease of reference, we will call the Level 1 claims in Group A, A1 claims, the Level 3 claims in Group B, B3 claims, and so on. We will thus assume that relevance extends two levels, such that the paraplegia and lost limb claims are relevant to the death claim, but the lost hand claims are not. Again, I shall also stipulate that two claims match with exactly one claim on the level above, and thus four claims will match with exactly one claim two levels above.¹¹ Lastly, to indicate that the First Horn Case is split into two stages: Stage 1 claims are those claims without brackets, and the claims added at Stage 2 are the claims in brackets.¹²

¹⁰ If one does not find this intuitive, consider lost fingers, broken arms, or lost toes or something similarly weak that is not relevant to the loss of a life.

¹¹ Of course, in reality no claims will match exactly with other claims like this, and we will end up with partial leftovers. For instance, a death claim might be loosely equal to one and a half claims. For simplicity, throughout the thesis I will treat claims as if they match exactly just to keep the cases manageable. However, this assumption does hide an ambiguity. I address this ambiguity in Appendix 11.

¹² This approach is borrowed from Aart van Gils and Patrick Tomlin (2019).

With this in view, we can see that at Stage 1, the thousand B4 claims are irrelevant to the A1 claim and thus Group A wins. However, at Stage 2 where a hundred Level 3 claims are added to Group A and six Level 3 claims are added to Group B, depending on which relevance approach we take and how we match the claims we might end saving Group B.

For instance, suppose we adopt a Local Relevance approach to limited aggregation. Suppose that this Local Relevance approach tells us to match the A1 claim with the B3 claims first.¹³ If we match the claims like this then the B3 claims will outweigh the A1 claim, and the B4 claims will outweigh the A3 claims. Thus, the view would tell us we ought to save Group B at Stage 2. Yet, this conclusion seems quite implausible. At Stage 2 we have strengthened Group A by more than Group B, yet it seems to have altered the decision of whom to save.

To make this problem clear, suppose that we are hospital directors choosing which ward to send our limited resources to. At first, Ward A has a dying man in it, and Ward B has a thousand people threatened by the loss of a hand. Because the loss of a hand is not relevant to the loss of a life, we choose to send our resources to Ward A. But just before we send the resources to Ward A, six people threatened by the loss of a limb are wheeled into Ward B and a hundred are wheeled into Ward A. But rather than say we have even more reason to send our resources to Ward A, we in fact swap to sending our resources to Ward B. This seems absurd!

In fact, we can sharpen the horn further than Horton suggests. Consider *First Horn Sharpened*:

First Horn Sharpened:

Stage 1 - You can save Group A or Group B. Group A contains a death claim. Group B contains one thousand lost hand claims.

¹³ Van Gils and Tomlin’s *Match to the Strongest Competing Claim* is such a view.

First Horn

Stage 2 - One hundred paraplegia claims are added to Group A. Six lost limb claims are added to Group B.

First Horn Sharpened

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia	(100)	
3 - Lost Limb		(6)
4 - Lost Hand		1000

In this case Group A is strengthened not just by more claims than Group B but also by stronger claims. Yet at Stage 2 certain limited aggregation accounts will still swap to Group B: the A1 claim will be outweighed by the B3 claims, and the A2 claims will be outweighed by the B4 claims.

To put this in terms of our example, we are in the process of sending the resources to Ward A, to save the dying man, when we find out that six people threatened by the loss of a limb have been wheeled into Ward B, and a hundred people threatened with paraplegia have been wheeled into Ward A. Even though Ward A has more, and stronger claims added to it, such a view will still tell us to swap to resourcing Ward B. This consequence is terrible.

Any view that lands on the First Horn of Horton's Dilemma thus violates two highly intuitive principles. In the *First Horn Case* such views violate:

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. (Van Gils & Tomlin, 2019, p. 253)

And in *First Horn Sharpened* they violate a new principle I call:

Greater Consideration for Stronger Claims: adding claims of differential strength and equal numbers cannot make the group to which stronger claims are added less choice-worthy compared with the group to which weaker claims are added.

Chapter 1 – Horton’s Dilemma

These seem like highly plausible principles. Yet any view that lands on the First Horn violates them both! Thus, I agree with Horton on the First Horn - it is devastating for any limited aggregation view that lands on it.

Now we might think that the First Horn is only a problem for Local Relevance views. However, Global Relevance views will also land on the First Horn and violate the *Principle of Net Addition* and *Greater Consideration for Stronger Claims* in different cases. Consider the following case:

First Horn for Global Relevance

Level	Group A	Group B
1 - Death		(2)
2 - Paraplegia	(1)	
3 - Lost Limb	10	
4 - Lost Hand		
5 - Lost Finger		100

In this case, at Stage 1 the B5 claims outweigh the A3 claims and so we should save Group B. But according to Global Relevance at Stage 2, when the death and paraplegia claims are added, we should swap, and save Group A instead. This is because the lost fingers are irrelevant compared to the paraplegia claim, and so are completely dismissed from consideration. This then allows the paraplegia and lost limb claims to outweigh the death claim. But of course, just like *First Horn Sharpened*, this means we have strengthened Group B by more and stronger claims and still swapped to saving Group A.

The First Horn then poses a serious problem to both types of Relevance view. However, there is an upshot of all this, if it will help us to narrow down the range of acceptable limited aggregation views. Any views which land on the First Horn and violate these principles should not be considered. Importantly this means that all Global Relevance views should be dismissed from consideration, as all Global Relevance views will land on this horn. Not all Local Relevance views land on this Horn, but a large number do, and these views should also be dismissed from consideration.

Second Horn

Given the large number of ways in which limited aggregation may be expressed, and the equally large number of different answers in different cases that these views will give, progress like this should not be scoffed at. If it engenders greater agreement among limited aggregationists by ruling out Global Relevance¹⁴ and brings us closer to understanding the structure of aggregation and the right way of weighing these claims, then this is truly very helpful. Of course, this upshot is reliant upon the Second Horn not being equally devastating, and that is what I will turn to next.

Second Horn

I will explain the Second Horn in my own terms, again for consistency of style. Consider the following two cases:

Case 1 - You can save Group A or Group B. Group A contains a death claim. Group B contains six lost limb claims.

Case 2 - You can save Group C or Group D. Group C contains one hundred lost limb claims. Group D contains one thousand lost hand claims.

Case 1

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia		
3 - Lost Limb		6
4 - Lost Hand		

Case 2

Level	Group C	Group D
1 - Death		
2 - Paraplegia		
3 - Lost Limb	100	
4 - Lost Hand		1000

¹⁴ The First Horn is not the only reason we have for dismissing Global Relevance. Patrick Tomlin provides some further devastating objections to the Global Relevance view in heterogeneous cases (Tomlin, *On Limited Aggregation*, 2017). We will return to Global Relevance in Chapter 3, where I will show how Global Relevance and Local Relevance can be combined in a way that does not cause the problems identified here.

Chapter 1 – Horton’s Dilemma

In Case 1, we should save Group B, and in Case 2, we should save Group D. Let us use our example to make the case clear. In Case 1, we choose to send our resources to Ward B to save the six people from losing their limbs, instead of saving the life of the man on Ward A. In Case 2, we choose to send our resources to Ward D to save a thousand lost hands, instead of sending our resources to the hundred people losing their limbs on Ward C.

Now let us consider the Second Horn Case, which is a combination of Cases 1 and 2. We can save Groups A and C or Groups B and D:

Second Horn

Level	Group A + C	Group B + D
1 - Death	1	
2 - Paraplegia		
3 - Lost Limb	100	6
4 - Lost Hand		1000

To put this in terms of our example, let us say that we realise that Wards A and C are merged by hospital administrators as are B and D. The hundred people in Ward C losing limbs are moved in with the dying man on Ward A, and the thousand people in Ward D losing hands are moved in with the six people losing limbs in Ward B, to form very crowded wards. Thus, if we save those on Ward A, we also save those on Ward C, and if we save those on Ward B, then we also save those on Ward D.

Now whom should we save in the Second Horn Case? Any Local Relevance view that wants to avoid the First Horn of the dilemma must save Group A+C. This is because the Second Horn Case is identical to the First Horn Case at Stage 2. Thus, if we are to avoid violating the *Principle of Net Addition and Greater Consideration for Stronger Claims*, we must save Group A+C.¹⁵

¹⁵ Global Relevance views will also tell us to save Group A+C, because the lost hand claims become irrelevant when the wards are combined. Importantly, this means Global Relevance views are skewered by both Horns of the dilemma.

Second Horn

But this seems very odd. When the pairwise comparisons were treated separately, in Cases 1 and 2, Group B and D were the winning groups. But when we combine the two winning groups and the two losing groups, without adding any more morally relevant features, we have swapped to saving the losing groups, A and C.

In this way, limited aggregation views which land on the Second Horn violate another highly intuitive principle, *Weak Additivity*.

Weak Additivity: If A is preferable to B, and C is preferable to D, then A with C is preferable to B with D.

To make this problem clear, consider our example again. In their individual pairings we would send our resources to Ward B, and save the six lost limbs instead of the person's life on Ward A. Similarly, we would send resources to Ward D and save a thousand lost hands instead of the hundred lost limbs on Ward C. So, Wards B and D are our winning wards and A and C are our losing wards.

But when Wards A and C, and B and D are merged by hospital administrators, we ought to send our resources to Ward A + C. We have chosen to save the losing patients over the winning patients from the previous match ups, without changing the strength or number of claims; we turn the two losing groups into the winning groups! This is highly counter-intuitive: it does not seem that combining claims in this way should change who wins.

The Second Horn is Unavoidable

Before we consider why the Second Horn is not as devastating as it might first seem, we should first strengthen the problem. Whilst Horton indicates that we could avoid the Second Horn by landing on the First Horn, this is not in fact true. The Second Horn is a bigger problem than perhaps first assumed - it is completely unavoidable for all limited aggregation views.

To see why consider Cases 3 and 4:

Chapter 1 – Horton’s Dilemma

Case 3 - You can save Group A or Group D. Group A contains a death claim. Group D contains one thousand lost hand claims.

Case 4 - You can save Group C or Group B. Group C contains one hundred lost limb claims. Group B contains six lost limb claims.

Case 3

Level	Group A	Group D
1 - Death	1	
2 - Paraplegia		
3 - Lost Limb		
4 - Lost Hand		1000

Case 4

Level	Group C	Group B
1 - Death		
2 - Paraplegia		
3 - Lost Limb	100	6
4 - Lost Hand		

In Case 3, limited aggregation views will tell us that we ought to save Group A - the lost hand claims are not relevant to the death claim. In Case 4, such views will tell us to save Group C. So, Group A and Group C are our winning groups.

But, of course, Cases 3 and 4 can be combined to form the Second Horn Case again. I think it should be clear now what the issue is. When split like this A and C are our winning groups, and B and D are our losing groups. But this is the opposite to what we find in Cases 1 and 2. In different comparisons the very same claims can be considered the winning claims or the losing claims!

Hence, limited aggregation views cannot avoid the Second Horn, and in one way or another must violate *Weak Additivity*.¹⁶ Therefore, in so far as there is no choice but to land on the Second Horn, it

¹⁶ Bastian Steuwer also recognises the unavoidability of the Second Horn (Steuwer, 2021a, pp. 31-34). He also attempts to show why this Second Horn is not devastating. Whilst he recognises that *Weak Additivity* (in his terms the ‘Principle of Agglomeration’) is the central issue, he fails to fully explain why and when the principle

Second Horn

is more prudent for a limited aggregationist to avoid the First Horn of the dilemma. In fact, the above case shows that Horton's dilemma is not really a dilemma at all, with the First Horn avoidable and the Second Horn not. Thus, limited aggregationist should not be concerned with the First Horn at all. We can avoid the First Horn without any downsides. Even if the Second Horn remains a problem, one problem is better than two.

The Second Horn is Not Devastating

So, we have seen both that the Second Horn violates *Weak Additivity*, and that all limited aggregation views, necessarily, land on the Second Horn. The obvious next step, therefore, is to consider how problematic it is to reject *Weak Additivity* and embrace the Second Horn. If doing so is unconscionable, then limited aggregation is sunk. On the other hand, if we can reject *Weak Additivity*, then we are going to want explanations for why and when we can.

Weak Additivity

So, let us consider how plausible the fundamental presumption of *Weak Additivity* really is.¹⁷ For starters, there are many day-to-day cases which violate *Weak Additivity*, such as complimentary and substitute goods. Such goods violate *Weak Additivity* because there is some important relation between the relevant values of each good. For instance, a book end is more valuable with its pair and a DVD film is less valuable if you already have a Blu-Ray copy. In principle then, we should not be surprised if there are ethical cases which violate *Weak Additivity* too.

fails and crucially when the principle does not fail. He mentions that there are interactions between the different claims but goes no further in his analysis. In doing so he fails to recognise the structural differences between the individual pairings and the combined case, and the role of dominance in determining the limit of appealing to interaction effects. He also neglects to see the core issue with the First Horn, and so misses that his own view also lands on a version of the First Horn (see Chapter 2). What follows is a more detailed and comprehensive approach to the dilemma, though one which shares many of Steuwer's insights. Thanks are owed to an anonymous reviewer for bringing this paper to my attention.

¹⁷ Much of what I say here mirrors what others say regarding transitivity in these cases. This ground is well enough covered by them and the much wider literature on transitivity, so I will keep my criticism specifically focused on *Weak Additivity*. For more on transitivity see Temkin (2012), Kamm (2007, pp. 484-486), Voorhoeve (2013) and Rachels (1998).

To make this even clearer, suppose you and I are picking team-mates for our weekly football match. By the time we are finished picking, I tactlessly start gloating, comparing my players to yours. I point out that my striker is better than yours, my goalkeeper is better than yours, and so on. In fact, for each match up, my player is stronger than yours. We then proceed to play, and your team wins comfortably. It turns out whilst I had been picking players by their individual strengths, you had been constructing a team that plays well together. Of course, this should be familiar, especially given that such cases form the basis of many sporting triumphs and the plot of almost every sports film.

Nevertheless, one might object by arguing that these cases are not analogous. In the football case there are obvious interaction effects between players, yet *Weak Additivity* only holds when we ignore interaction effects. But in the same way that there are interaction effects in the football case, limited aggregationists argue that there are interaction effects in the aggregation cases.

Limited aggregationists will appeal to interaction effects in our case by arguing that claims against dying alter the relevance of claims against lost hands. For instance, my claim against a lost hand becomes less pressing, and demands less attention, in the light of your claim against death.

Depending on the account, we might say that it would be disrespectful to the dying person to take my claim into consideration (Kamm, 2007, p. Chapter 10); or that my claim’s duty grounding force is defeated by the presence of the death claim (Tadros, 2019); or that it would not be permissible for me to prefer my own hand to be saved over your life (Voorhoeve, 2014); or that my claim becomes less sympathetic in light of the death claim (Voorhoeve, 2017). In this way, we should not be surprised if limited aggregation views do violate *Weak Additivity*. In fact, given such interaction effects, it would be considerably more unusual if they did not, on occasion, violate *Weak Additivity*.

Now I need to be clear about what I am and am not saying here. My position here is not that *Weak Additivity* always fails, but that sometimes, under certain conditions, it fails. Not all violations of

Second Horn

Weak Additivity are valid¹⁸ and we need good explanations if we are to accept any violations of *Weak Additivity*. In this way we might consider *Weak Additivity* to be a default assumption - we assume *Weak Additivity* other than where we have reason to doubt it. For instance, we have good reasons to reject *Weak Additivity* in the bookend, film, and football cases; the first bookend is more useful with the second, the DVD is less valuable because we already own the film, your football team is better because your players understand each other, etc.

Therefore, if we are to accept such violations in the aggregation cases, we will need to give good reasons in each case as to why *Weak Additivity* is violated. We will need to specify how exactly the interaction effects in our cases lead to the violation of *Weak Additivity*; if not, then no appeal to the failure of *Weak Additivity* in the general will suffice.

Thus, let us return to our hospital case. Before we amalgamate the wards, we choose to save the six lost limb claimants in Ward B instead of the death claimant in Ward A. We also choose to save the thousand lost hand claimants in Ward D instead of the hundred lost limb claimants in Ward C. What then explains the reason we ought to save Ward A+C when the wards are combined as per the Second Horn Case?

Second Horn

Level	Group A + C	Group B + D
1 - Death	1	
2 - Paraplegia		
3 - Lost Limb	100	6
4 - Lost Hand		1000

The first thing to note is that when the wards are amalgamated each claim comes into competition with new claims. At first the lost hand claims are only in competition with the lost limb claims, now they are in competition with a death claim as well as the lost limb claims. So, the lost hand claims are

¹⁸The *Principle of Net Addition* and *Greater Consideration for Stronger Claims* are such examples. See Dominance section below.

now in competition with claims they are relevant to and a claim they are not relevant to. Now, according to Local Relevance views, the lost hand claims might still be relevant to what we ought to do on the whole, but only if the death claim is outweighed by other considerations. If the death claim is not outweighed by other considerations, then the lost hands are irrelevant and can play no role in our decision.

But we might think this is no explanation of why we ought to swap to Ward A+C when they are amalgamated. In this case, it seems that the death claim is outweighed by other claims, namely the lost limb claims in Ward B. So, the lost hand claims remain relevant, outweigh the lost limb claims in Ward C, and so we ought to save Ward B+D.

However, there is a second interaction effect that needs to be considered. The lost hand claims are not the only claims to have had changed what they are in competition with. The other claims also come into competition with different strength claims, and this might further change whom we ought to save. In this case, before amalgamation, the six lost limb claimants in Ward B were only in competition with the death claimant. After amalgamation they are now also in competition with the lost limb claimants in Ward C.

This change is important, because, depending on how we compare the claims, the lost limb claims might not be compared to and outweigh the death claims. If the death claims are not outweighed, then the lost hand claims are not relevant and cannot be taken into consideration. Thus, if we say that, in this case, the lost limb claims are outweighed by the other lost limb claims, then the death claim is not outweighed. So, the question comes down to how we ought to match claims when the wards are amalgamated, i.e., what is the best matching principle?

This is a bigger debate than we have time to go into in this chapter, but there may be a number of different principles or reasons that might explain why we ought to match the lost limb claims with the other lost limb claims at Stage 2. For example, a matching principle from the literature that

Second Horn

might explain this is *Strongest Decides* (Van Gils & Tomlin, 2019). This principle tells us that the person with the strongest claim has the power to decide how claims ought to be matched. In brief, the argument is that the strongest claimant has the most to lose, and so (in keeping with the anti-aggregative parts of limited aggregation) should have greater influence over whom we save or, at least, tiebreak between different ways of matching claims. *Strongest Decides* would tell us at Stage 2 to match the lost limb claims with the other lost limb claims (because it benefits the strongest overall claimant) and thus we ought to save Group A+C at Stage 2.

Principles such as *Strongest Decides* would be able to explain why the interaction effects introduced when the groups are amalgamated mean we ought to swap whom we save in the Second Horn Case. As such, they will also be able to explain why *Weak Additivity* fails in these cases. Now, we will return to the question of matching principles in the next chapter, where I will explain why I think *Strongest Decides*, and all other so far identified matching principles, fail. However, I will develop a new matching principle that can explain why we ought to match the lost limb claims with the paraplegia claims in the Second Horn Case and thus explain why we ought to save Ward A+C. Hopefully, what we have here, however, is sufficient to demonstrate that violations of *Weak Additivity* are theoretically justifiable.

Horton's example

Lastly, we ought to look at Horton's argument against embracing the Second Horn. Horton provides an example which purports to show that embracing the Second Horn is terrible:

Suppose that on your left are two buttons marked [A]¹⁹ and [B], and on your right are two buttons marked [C] and [D]. If you press a button, that will save the corresponding group. But you can press only one button in each pair. Suppose next that your arm span is just slightly too short for you to reach both pairs of buttons simultaneously. It follows, on the

¹⁹ Button letters changed to match the Second Horn Case.

view that we are considering, that you should press button [B] on your left and then button [D] on your right, even though, had your arm span been just slightly longer, it would have been permissible for you to simultaneously press buttons [A] and [C]. (Horton, 2018, p. 173)

Horton’s illustration seems to suggest that if your arm span is too short to press the buttons simultaneously, then you are dealing with Cases 1 and 2 separately. But if your arm span is long enough you are dealing with the Second Horn Case. Given the answer we get in the Second Horn Case is different to the answer we get in Cases 1 and 2, whom you save is implausibly dependent on your arm span, despite no other features of the case changing.²⁰

This example does seem devastating at first glance. However, Horton misapplies his example. It does not provide pressure against violating *Weak Additivity* like Horton supposes, but instead puts pressure on when we should treat decisions as two separate decisions or one larger decision.

Thus, whether we treat the decision between A and B and the decision between C and D as one decision is not about whether the buttons are pressed simultaneously, but whether we have true free choice across the groups or whether there are interaction effects. If there are interaction effects between the cases, such as restricted choice across the groups, then we should treat the decision as one. If not, then we should treat the cases separately.

In this way, Horton’s example fails to capture an important interaction effect of the Second Horn Case. In his example, the choices are completely independent of each other: we could choose to save Group B in the first competition and then still have an open choice between saving Group C or D. But in the Second Horn Case – where the choice is between Group A+C and Group B+D, we can only choose to save Group B if we also choose to save Group D. By making a choice in one of the competitions we constrain whom we can save in the other competition.

²⁰ This is not a case of ought implies can. I can save any of the claimants, the only thing that changes with my arm span is whether I can save people simultaneously.

Second Horn

So, we must amend Horton's example if it is to accurately capture the Second Horn Case. Suppose, then, that if I choose to press button B on my left, I cannot press button C on my right – let us say both buttons use the same wire, and two signals cannot be sent down the same wire. Similarly, if I choose to press button D on my right then I cannot press button A on my left, and vice-versa.

With a more accurate illustration in view, we can see that relevance is going to be important once again. Suppose I choose to address the buttons on my right first. If I press button D and save a thousand hands, then I can no longer press button A on my left and save the dying man. Here, by choosing to save a thousand hands I leave a person to die. Of course, this is impermissible on any limited aggregation view.

Similarly, suppose I go to press button B on my left first. If I press button B and save six limbs, then I can no longer press button C on my right and save one hundred limbs. Of course, this too would be impermissible. Thus, there is no way to handle these cases independently. Making a choice in one case changes the available options in the other case, and so the only way we can handle the situation is to treat both choices together.

Therefore, neither the length of one's reach nor whether the buttons are pressed simultaneously has any effect on whether the context changes. We should treat the choices separately if there are no interaction effects, even if we can press the buttons simultaneously. Correspondingly, we should treat the choices as one if there are sufficiently strong interaction effects, even if we cannot press the buttons at the same time. Horton's illustration therefore fails to show what it purports to.

Dominance

One last question remains however: if we are going to allow that combining two losing groups can turn them into the winning group, then why do we not go the whole-hog and abandon the *Principle of Net Addition* and *Greater Consideration for Stronger Claims*? Why is the First Horn devastating but the Second Horn not? It certainly seems that the problem with the First Horn cases is that the

addition of claims should not make the group to which the losing claims are added more choice-worthy. If we are to maintain that the First Horn should be avoided, then we need a reason to treat violations of the *Principle of Net Addition* and *Greater Consideration for Stronger Claims* differently to violations of the Second Horn.

Firstly, I think it is important to reiterate that not all violations of *Weak Additivity* are valid. As I have argued above, we should accept that, by default, *Weak Additivity* applies in any given case - we need good reason to abandon it. Furthermore, whilst the *Principle of Net Addition* and *Greater Consideration for Stronger Claims* are more specific versions of *Weak Additivity*, they are not principles we should want to reject; they are independently highly intuitive. It is important to note that whilst rejecting the *Principle of Net Addition* or *Greater Consideration for Stronger Claims* entails rejecting *Weak Additivity*, rejecting *Weak Additivity* does not entail rejecting the *Principle of Net Addition* or *Greater Consideration for Stronger Claims*. If the engine of my car is broken, it entails that my car is broken, but my car being broken does not entail that the engine is broken.

Thus, my argument against *Weak Additivity* should not be taken as a sufficient reason to reject the *Principle of Net Addition* or *Greater Consideration for Stronger Claims* by itself - if one wants to reject these principles one needs separate reasons to do so. As it stands, I see no particularly good reason to reject them, and if an account can be developed which maintains the principles, it will have one less bullet to bite. I develop such an account in the Chapter3.

Secondly, I think we have good theoretical reason to want to distinguish between *Weak Additivity* and the *Principle of Net Addition* and *Greater Consideration for Stronger Claims*, and good reason not to violate the latter two principles. The reason lies in the notion of dominance:

Dominance: A claim, or set of claims, dominates another claim, or set of claims, if it is better in some respects and at least equal in all other respects.

Second Horn

We can see that in the cases that violate the *Principle of Net Addition* and *Greater Consideration for Stronger Claims* the stronger claims dominate the weaker claims. They are at least equal in all respects and stronger in at least one. For instance, in the *First Horn Case* we add equal strength lost limb claims to both groups, but more claims to Group B. Thus, the claims added to Group B dominate the claims added to Group A. So views that land on the First Horn violate a principle I call:

Principle of Added Dominance: adding claims cannot make the group to which dominant claims are added less choice-worthy compared with the group to which the dominated claims are added.

But why does dominance matter? What is wrong with violating the *Principle of Added Dominance*? The answer is that if we violate the *Principle of Added Dominance*, we are ignoring some morally relevant features when all other features are equal. One group is made stronger in one respect and not made weaker in any respect - yet we still swap to saving the other group.

Compare this to the Second Horn cases which do not violate this *Principle of Added Dominance*. In the Second Horn cases we add claims that are stronger than each other in different ways. One set of claims is greater in number, the other set of claims has stronger individual claims. So, in the Second Horn cases, the winning claims do not dominate the losing claims. They might be on balance stronger, but they are still weaker in one dimension.

This explains why we can save the two losing groups in the Second Horn cases. According to limited aggregation views, whether the numbers matter depends on the other claims in the competition. Weak claims cannot aggregate their large numbers if they are in competition with much stronger claims. Thus, when the groups are combined in the Second Horn cases, we need to reconsider whether the numbers matter. Claims that were previously irrelevant in a two-way tie might now find themselves relevant again when the claims they were irrelevant to are matched by other claims. Similarly, claims that were relevant in their individual pairings might now come into competition with significantly stronger unmatched claims and thus lose relevance.

Chapter 1 – Horton’s Dilemma

To see this, consider the Second Horn Case again, but this time split into two stages:

Second Horn Case:

Stage 1 - You can save Group A or Group B. Group A contains a death claim. Group B contains six lost limb claims.

Stage 2 - Group C is added to Group A, Group D is added to Group B. Group C contains one hundred lost limb claims. Group D contains one thousand lost hand claims.

Second Horn Case

Level	Group A + C	Group B + D
1 - Death	1	
2 - Paraplegia		
3 - Lost Limb	(100)	6
4 - Lost Hand		(1000)

At the first stage Group B wins. At the second stage Group B is strengthened by the winning Group D claims, whilst Group A is strengthened by the losing Group C claims. However, Group B is strengthened by a greater number of claims, whilst Group A is strengthened by stronger claims. We can see that the additional lost hand claims are now irrelevant to the strongest claim in the competition. Thus, it should not be surprising that Group A+C becomes the winning group. Group A has been strengthened by claims that are relevant in all match ups, whilst Group B has been strengthened by claims that are irrelevant in at least one match up.

Summary

Horton’s dilemma then should not be considered fatal to all limited aggregationists. Firstly, I have shown that it is not really a dilemma, with the First Horn avoidable and the Second Horn not. The First Horn then is a problem only for a certain range of limited aggregation accounts. Given the severity of the First Horn and the unavoidability of the Second Horn, I suggest all limited aggregationists abandon accounts that land on the First Horn of the dilemma. Even if one is not

Summary

convinced by my argument that the Second Horn is not a problem for limited aggregation views, two Horns are worse than one. Thus, at the very least, this chapter shows that limited aggregationists should focus on the Second Horn.

Secondly, once we recognise that the Second Horn rests on an assumption of *Weak Additivity*, we can see exactly why limited aggregationists should not be concerned with the Second Horn. *Weak Additivity* is an implausible assumption for any limited aggregationist to hold. Where limited aggregationists maintain that numbers sometimes count and sometimes do not, that sometimes it is right to add claims and sometimes it is not, then the very core of limited aggregation is to reject claims about additivity. Furthermore, I have shown that the counter-example that Horton provides against the Second Horn is faulty: it does not capture all the important features of the Second Horn cases. Once we amend his example to properly capture these features then the limited aggregation approach seems not just plausible but in fact quite intuitive.

Thus, I have shown that *Weak Additivity* is not only an implausible assumption for limited aggregationists to hold in principle, but I have also identified exactly where it goes wrong in these Second Horn cases. Particular interaction effects indicate a deeper justification to the rejection - namely that more appropriate comparisons arise when the groups are amalgamated. Thus, Horton fails to establish that limited aggregation views, necessarily, have implausible consequences, either in principle or in practice.

Furthermore, I have also begun to outline the conditions under which we can expect *Weak Additivity* to fail, and when it ought to be maintained. The concept of dominance allows us to distinguish between cases where *Weak Additivity* might fail and cases where it definitely ought not to. It also begins to give us an insight into the reasons why we might think that *Weak Additivity* fails.

Nevertheless, this is only a partial picture. To fully understand the deeper justification for why we ought to swap whom we save in the Second Horn cases, we need to know exactly how we ought to match claims in these heterogenous cases. I will turn to this issue in the next two chapters.

Chapter 1 – Horton’s Dilemma

Hopefully, however, this chapter shall put to bed Horton’s criticism for good and allow limited aggregationists to focus their energies on developing the most plausible versions of their views.²¹ Nevertheless, Horton’s dilemma does bring to the fore the issues contained in the First Horn and shines a light on a number of principles any plausible limited aggregation view will need to maintain. As such, it does help us to realise the range of acceptable views. In the next chapter, we will look at a range of views from the literature and show that none of them meet this standard. Then in Chapter 3 I will outline a view that does meet these standards and can complete the explanation for why the Second Horn is not a problem.

²¹ Horton does actually provide another objection to limited aggregation views, based on cases involving risk (Horton, 2020), which we might think we ought to address before moving on. However, this argument is weaker, is limited to more controversial cases of risk, and shares the same underlying problematic assumptions as the ‘fatal dilemma’ considered here. Namely, the argument once again assumes *Weak Additivity* in cases where there are clearly interaction effects. As such (and because my thesis will not focus on cases involving risk) I will not address this problem in the main thesis. However, Appendix 1 briefly explains my objections to this problem.

Chapter 2 – Local Relevance Views and their Matching Principles

Introduction

Local Relevance views have significant appeal. They successfully capture our intuitions in aggregation cases, including in the heterogenous cases that previous relevance views struggled with. They do so by matching or balancing competing claims, until one group has claims which can no longer be matched. In this chapter I consider the different ways we might match or balance competing claims. This is a question that has not been sufficiently addressed in the literature, yet Local Relevance views live or die by their matching principles. I show that all the matching principles so far offered by relevance theorists fail. They violate several commonly held principles and break the internal logic of relevance views.

In the previous chapter we introduced the distinction between Global and Local Relevance. I showed there that Global Relevance is skewered by both Horns of Horton's dilemma, and that any view that lands on the First Horn of the dilemma should be dismissed from consideration. Thus, we shall set aside Global Relevance in this chapter, though we will return to it in Chapter 4.

As we saw in the last chapter Local Relevance views are designed to cope with heterogenous cases i.e., cases where each group contains claims of differing strength. Remember that the core idea to Local Relevance is that claims are not relevant or irrelevant simpliciter. Instead, a claim might be relevant compared to some claims in the competition but irrelevant compared to others. This means that a claim might be irrelevant compared to a stronger claim but remain relevant to a weaker claim and thus relevant to our overall decision. If the stronger claim is matched (or counterbalanced) by some other claim, then the weaker claim can play a role in deciding which group to save.

Now, because Local Relevance views treat the relevance of claims in this way, they admit of a number of different interpretations. This also means that different Local Relevance views might give different answers in different cases. To make this clear consider the follow case:

Case 5

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia	2	2
3 - Lost Limb		
4 - Lost Hand		100

Remember that for illustrative purposes we have stipulated that claims are relevant to competing claims if they are within two levels of the claim they compete with. So, in this case the paraplegia claims are relevant to the death claim, but the lost hand claims are not. We also stipulated that one claim is matched by exactly two claims on the level below, such that two paraplegia claims would balance with a single death claim.

In Case 5 then, we can see that Local Relevance could be formulated such that Group A or Group B wins. If the death claim is counterbalanced by the paraplegia claims in Group B, then the lost hand claims can play a role in whom we ought to save. The lost hand claims would outweigh the paraplegia claims in Group A and so we ought to save Group B. However, if the paraplegia claims in Group B are outweighed by the paraplegia claims in Group A, then the death claim is not counterbalanced. This means that the lost hand claims remain irrelevant and so we ought to save Group A.

Thus, much rides on the account of local relevance we take, specifically, on how we match the claims in complex heterogenous cases. In this chapter, I want to consider the different options put forward for how we might match or balance competing claims. I shall show that all the matching principles, so far offered, fail. They violate several commonly held principles and break the internal logic of relevance views.

Introduction

These views are *Strongest Decides* and *Sequential Claims Matching* provided by Van Gils and Tomlin, and Bastian Steuwer's *Hybrid Balance Relevant Claims*. We will then look at the operating principle behind the latter two views, *Match in Order of Strength*. We will consider what this principle gets right and wrong and how it causes problems for the aforementioned views. Lastly, we will look at an alternative principle, *Match in Reverse Order of Strength*, and why this principle also fails.

I shall go beyond just showing how the views fail, as I will diagnose why they get the wrong answers. Doing so will allow us to identify what these views get right and what they get wrong, and will allow us to take forward the helpful elements of these views. This will set the groundwork for the next chapter in which I will provide my own matching principle, which will avoid the problems identified here.

Two Aspects of Matching

Before we turn to the particular views it is important to note that the question of how to match claims is actually made up of two sub-questions. Both of which will need answering to give a full account of how to match claims without ambiguity. These are: (1) in what order do we match and set aside claims? and (2) which claims do we match each particular claim to?

To see why this is important, consider again Case 5:

Case 5

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia	2	2
3 - Lost Limb		
4 - Lost Hand		100

Firstly, we have the question of which order we ought to match the claims. If we match the A1 claim first, it will be matched by the B2 claims, and the B4 claims will then outweigh the A2 claims. Thus, we will end up saving Group B. However, if we instead decide to match the B2 claims first (and we

decide to match the B2 claims with the A2 claims), then we will save Group A because the B4 claims remain irrelevant to the A1 claim.

Secondly, we have the question of which claims we match to which. In this case, if we decide to match the B2 claims first, there are actually two groups of relevant claims we might match the B2 claims to. If we match them to the A2 claims, then Group A wins. But if we match them to the A1 claim, then Group B wins. Thus, both the order of matching and which claims are matched to which will affect whom we ought to save.

Strongest Decides

With this distinction in view, let us now turn to the views from the literature. Let us start with the simplest of these views - *Strongest Decides*. Van Gils and Tomlin identify this principle as a potential alternative to their preferred Sequential Claims Matching view (Van Gils & Tomlin, 2019, pp. 250-252). The following is a simple formulation of the principle:

Strongest Decides: Claims should be matched in whichever order and whatever way is in the interest of strongest overall claimant.

In essence, *Strongest Decides* tells us that if it is possible to match in a way to save the claimant with the most at stake, then we ought to do so. In other words, the claimant with the most at stake gets to choose how to match the claims. In this way, *Strongest Decides* is quite simple as it gives the same answer to both sub-questions - we defer to the strongest claimant about both the order of matching and which claims are matched with which. In Case 5, this means matching the B2 claims first, and matching them to the A2 claims, such that Group A wins.

We can see then that *Strongest Decides* gives very strong priority to the worst-off and is on the anti-aggregative end of the spectrum of limited aggregation views. Thus, for those with strong anti-aggregative leanings *Strongest Decides* might be quite appealing. Nevertheless, I think *Strongest Decides* fails.

Strongest Decides

Firstly, I have a couple of theoretical worries. It strikes me that *Strongest Decides* gives too much priority to the strongest claimant. The strongest claimant seems to have the power to gerrymander the way claims are matched such that the right conclusion for her (i.e., the conclusion in which she is saved) is the one we arrive at. Without a proper process decided on from the outset this seems unfair.

What *Strongest Decides* requires is also going to be ambiguous in a large number of cases: what do we do when there are equal numbers of the strongest claims in both groups? Perhaps we let the next strongest claim decide. But such a solution only passes the problem down one layer. Plus, it seems to get further away from the anti-aggregative justification behind *Strongest Decides*, namely by deciding in favour of the second worst-off person.

Secondly, even if *Strongest Decides* can handle these worries, there are some deeper issues.

Strongest Decides does not land on the fatal First Horn of Horton's dilemma in the way it is presented in the previous chapter. However, it does land on the horn in other ways, or at least violates the same principles that lie behind the First Horn. Firstly, Van Gils and Tomlin debate whether *Strongest Decides* violates a principle they call *Equal Consideration for Equal Claims* (Van Gils & Tomlin, 2019, p. 251):

Equal Consideration for Equal Claims: all claims of equal strength ought to be given equal weight in determining which group to save. (Tomlin, On Limited Aggregation., 2017, p. 243)

We can see why in the following case:

Case 6

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia		
3 - Lost Limb		5
4 - Lost Hand	(10)	(10)

Case 6 consists of two stages. At Stage 1 the case involves only the claims without brackets, i.e., the A1 and B3 claims. At Stage 2 new claims are added to the case, indicated by the claims with brackets.

So, what does *Strongest Decides*, tell us to do in this case? At Stage 1 there is only one way to match the claims, we must compare the death claim with the lost limb claims, and the lost limb claims outweigh the death claim so Group B wins. However, at Stage 2 there is an ambiguity about how to match the claims. We could match the lost limb claims with either the death claim or the lost hand claims. If we match the lost limb claims with the death claim, Group B wins. If we match the lost limb claims with the lost hand claims, Group A wins. *Strongest Decides*, of course, will tell us to match such that Group A wins, and so will match the lost limb claims with the lost hand claims.

Here is where the problem lies. It seems that by matching this way we are giving extra weight to the lost hand claims in Group A to those in Group B. The lost hand claims in Group A are taken into consideration and weighed against the lost limb claims. But the lost hand claims in Group B are compared with the death claim and so are dismissed as irrelevant and not considered at all. This seems wrong.

Van Gils and Tomlin conclude, however, that *Strongest Decides* does not in fact violate *Equal Consideration for Equal Claims*. According to Tomlin, *Equal Consideration for Equal Claims* does not entail that all claims must be treated equally, but rather just that they are accorded equal weight (Tomlin, *On Limited Aggregation.*, 2017, p. 11). Thus, they argue that what “*Equal Consideration for Equal Claims* [my italics] 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” (Van Gils & Tomlin, 2019, p. 251). They argue that *Strongest Decides* gives equal claims equal weight and does not rule out some views from the outset, and so does not violate *Equal Consideration for Equal Claims*.

However, I find this deeply unsatisfactory. Claimants are less concerned with being given equal consideration at the outset than with being given equal consideration at the end of deliberation. It is

Strongest Decides

little compensation to be told your claim was given equal weight at the outset, if by the end of deliberation your claim is treated as irrelevant whilst another's identical claim has been instrumental in deciding what to do.

Nevertheless, even if one disagrees with my assessment of *Equal Consideration for Equal Claims* here, *Strongest Decides* definitely violates two other important principles mentioned in the previous chapter:

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. (Van Gils & Tomlin, 2019, p. 253)

And the new principle I called:

Greater Consideration for Stronger Claims: adding claims of differential strength and equal numbers cannot make the group to which stronger claims are added less choice-worthy compared with the group to which weaker claims are added.

In order to see why, we need to consider another new case.

Case 7

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia		
3 - Lost Limb		5
4 - Lost Hand		(100)
5 - Lost Finger	(20)	

In this case at Stage 1 the B3 claims outweigh the A1 claim. At Stage 2, one hundred Level 4 claims are added to Group B. Simultaneously, twenty Level 5 claims are added to Group A. Significantly more, and stronger, claims are added to Group B, yet *Strongest Decides* requires us to swap whom we save to Group A. This is because it would be favourable to the overall strongest claimant to match the A5 claims with the B3 claims. After that, the remaining B4 claims could not compete with

the strongest claim. Thus, this matching principle violates the *Principle of Net Addition* and *Greater Consideration for Stronger Claims*. This result is devastating to this matching principle.

Sequential Claims Matching

With *Strongest Decides* set aside, let us now look at Van Gils and Tomlin's preferred approach *Sequential Claims Matching*. This view is rather more complicated than *Strongest Decides*, so rather than attempt to explain the view, I think it is most useful just to present it and walk slowly through a test case to make clear how it works in practice. The following is the sequence:

"Sequential Claims-Matching

- I. Identify the strongest claim-type T1. Does one group contain more individuals with claims of type T1 than the other?
 - a. If not, eliminate all T1-claims from consideration.
- II. If so, match each T1-claim from the group with fewer T1-claims to a T1-claim from the group with more T1 claims, and remove the matched claims from consideration.
- III. 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?
 - a. If not, you should decide in favour of the group with the remaining T1 claims.
- IV. If so, do the relevant competing claims outweigh the T1-claims?
 - a. If not, you should decide in favour of the group with the remaining T1 claims.
- V. If so, match the set of remaining T1 claims to a set of relevant competing claims with comparable weight, and remove the matched claims from consideration.
- VI. Now consider the remaining unmatched claims. Of these, identify the strongest claim-type T2. Repeat the above procedure.
- VII. Continue until either:
 - a. one group contains unmatched claims, in which case you should decide in favour of that group; or

Sequential Claims Matching

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).” (Van Gils & Tomlin, 2019, pp. 228-229)

There is an important terminological point to note here. Van Gils and Tomlin call the strongest unmatched claim the ‘anchoring claim’. For instance, in Case 8 below, the Level 1 claims are the anchoring claims, until they are all matched, and then the Level 3 claims become the new anchoring claims. We call them ‘anchoring claims’ because it is to these claims that we determine whether a competing claim is relevant.

Now to see how *Sequential Claims Matching* works in practice, consider the following complex test case. The complexity of this case makes it hard to handle without guidance.

Case 8

Level	Group A	Group B
1 - Death	2	1
2 - Paraplegia		
3 - Lost Limb		5
4 - Lost Hand	4	

In this case the B3 claims are relevant to the A1 claims, but the A4 claims are not relevant to the B1 claim. A lost limb claim is relevant to death, but a lost hand claim is not relevant. With this in view, let us walk through our test case to see how the sequence works.

Steps I and II: we identify the strongest claims - the Level 1 claims - as our anchoring claims. Given that there are Level 1 claims in both groups we match and dismiss claims from consideration until one group has no Level 1 claims remaining. To demonstrate this, I have struck them out and replaced the remaining number of Level 1 claims in square brackets:

Case 8

Level	Group A	Group B
1 - Death	2-[1]	±
2 - Paraplegia		
3 - Lost Limb		5
4 - Lost Hand	4	

Steps III, IV and V: We need to check if there are any claims relevant to the remaining anchoring claim - the A1 claim. Given that relevance extends two levels up, we can see that the B3 claims are relevant and can therefore match with the A1 claim. Given that exactly four claims match with one claim two levels above, we match four B3 claims with the remaining A1 claim, leaving one B3 claim:

Case 8

Level	Group A	Group B
1 - Death	2-[1]	±
2 - Paraplegia		
3 - Lost Limb		-5-[1]
4 - Lost Hand	4	

Step VI: we repeat the previous steps but with the remaining B3 claim as the new anchoring claim. As we can see there are four relevant A4 claims. As the A4 claims are a level below the B3 claim, two A4 claims are needed to match the B3 claim. Therefore, the four A4 claims outweigh the B3 claim, and as there are no further relevant claims, Group A wins. The view is quite simple and intuitive to follow once we see it in practice.

In essence, *Sequential Claims Matching* starts with the strongest claims and matches those claims with relevant claims taking them out of consideration. Once the strongest claims are matched the sequence tell us to match the next strongest unmatched claims in the same way, until no more matchings are possible.

Sequential Claims Matching has a little more going for it compared to *Strongest Decides*. Firstly, it seems to have a good justification. Van Gils and Tomlin argue that “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

Sequential Claims Matching

must be matched or defeated by claims relevant to it, seems like a natural extension of the limited aggregation approach” (Van Gils & Tomlin, 2019, p. 228). Nevertheless, I think that *Sequential Claims Matching* fails.

Firstly, we should note that *Sequential Claims Matching* only answers the first sub-question. It gives us an answer to what order we should match claims, but not to which claims we should match claims to. Thus, *Sequential Claims Matching* will be ambiguous in a number of cases and is an incomplete view.

Now, Van Gils and Tomlin do recognise this ambiguity and so offer a number of supplementary matching principles for these cases (Van Gils & Tomlin, 2019, pp. 246-247). In considering how to match claims they focus their attention on these supplementary principles, a move that I think was misguided. Whichever supplementary principles they choose, *Sequential Claims Matching* will still fail, because the underlying problem is with their answer to the first of these sub-questions - the problem lies with the order in which they match claims. As such I shall set aside these supplementary principles and focus on *Sequential Claims Matching* directly.

The problem with *Sequential Claims Matching* is that it lands on the devastating First Horn of Horton’s dilemma. Thus, like *Strongest Decides*, it violates the *Principle of Net Addition* and *Greater Consideration for Stronger Claims*. To see why let us reconsider *First Horn Sharpened* from the previous chapter:

First Horn Sharpened

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia	(100)	
3 - Lost Limb		(6)
4 - Lost Hand		1000

In *First Horn Sharpened* at Stage 1, we save Group A and the strongest claimant. At Stage 2 we add one-hundred paraplegia claims to Group A and six lost limb claims to Group B. Despite this

Sequential Claims Matching tells us to match the strongest claim first, and so we end up matching the A1 claim with the only relevant competing claims - the B3 claims. With the A1 claim set aside, the remaining B3 claims with the B4 claims outweigh the A2 claims and thus we swap to saving Group B, violating both the *Principle of Net Addition* and *Greater Consideration for Stronger Claims*. Again, these principles are sufficient to dismiss *Sequential Claims Matching*; at least as it stands.

Hybrid Balance Relevant Claims

Let us now turn to Bastian Steuwer's view:

Hybrid Balance Relevant Claims. Relevant individual claims ought to be balanced against one another, starting with the strongest claim(s) overall. If there are relevant claims on one side that are not balanced by claims on the other, then these unmatched claims will be decisive in determining what we ought to do. If the claims are evenly balanced, then we are permitted to save either group, or perhaps required to give equal chances. The relevance of claims is determined by two conditions:

(1) *The local relevance condition*: A claim can be balanced against another claim only if the two claims are relevant to one another.

(2) *The global relevance condition*: Every individual with a strong claim has a veto against the consideration of any type of claim that is irrelevant to her claim, if such consideration would lead to her claim not being satisfied (e.g., by not being saved). (Steuwer, 2021a, pp. 18-19)

As we can see here, the core of Steuwer's view is very similar to Van Gils and Tomlin's. It starts by matching the strongest claim and continues in a similar fashion. Therefore, we might think that the view fails in exactly the same way as *Sequential Claims Matching*.

However, Steuwer's version does actually manage to avoid the First Horn in the cases that pose a problem for *Sequential Claims Matching* above. It does so by introducing a veto for the strongest

Hybrid Balance Relevant Claims

claimants, via the global relevance condition. This veto means that claims which are matched and set aside can still play a role in determining whether a claim is relevant.

Thus, in *First Horn Sharpened*, above, the A1 claim maintains a veto against Level 4 claims. Thus, even when the A1 claims are matched and taken out of consideration, the B4 claims remain irrelevant. This, of course, means that Group A wins.

Nevertheless, Steuwer's view still fails. Firstly, his view is incomplete. It tells us only the order in which to match claims and not to which claims the strongest claim ought to be matched. Thus, this view is just as ambiguous as Van Gils and Tomlin's and will need supplementary principles to tell us what to do in those cases.

Secondly, I am sceptical of Steuwer's use of the veto. The veto seems to me too arbitrary and unfair to those weaker claims. Consider the following case:

Case 9

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia		2
3 - Lost Limb		
4 - Lost Hand	20	19 (+2)

At the first stage the lost hand claims matter. They are decisive in favour of Group A: the strongest claimant has decided that the lost hand claimants all deserve sympathy and should be included in the moral decision. But at the second stage when two extra lost hand claims are added to Group B, the strongest claimant does an about-face and decides that the lost hand claims no longer matter, leaving us with a tie. This seems very morally inconsistent, and I think the lost hand claimants would rightfully feel aggrieved by this change. In this case the lost hand claimants do not seem to be treated with respect or as ends in themselves, but merely as instrumentally useful to the strongest claimant.

Most problematically though, the veto does not even properly solve the issue. Steuwer’s view still violates both the *Principle of Net Addition* and *Greater Consideration for Stronger Claims* and so lands on the devastating First Horn of Horton’s dilemma. To see why, let us reconsider the following case:

First Horn for Global Relevance

Level	Group A	Group B
1 - Death		(2)
2 - Paraplegia	(1)	
3 - Lost Limb	10	
4 - Lost Hand		
5 - Lost Finger		100

In this case at Stage 1, the B5 claims outweigh the A3 claims and so Group B wins. At Stage 2, however, we add two death claims to Group B and one paraplegia claim to Group A. Thus, Group B is strengthened by more and stronger claims. Nevertheless, Steuwer’s view will still tell us to swap and save Group A. This is because the A2 claims have a veto against the B5 claims, and so the B5 claims are not taken into consideration. With the B5 claims dismissed, the Group A claims together outweigh the B1 claims. Thus, we can see that Steuwer’s view, because it incorporates Global Relevance, also violates both the *Principle of Net Addition* and *Greater Consideration for Stronger Claims*.

Now, there is an argument that the violations in this case are not as bad because the claims that are being added are the strongest claims. The argument is that when we add the strongest claims, we should expect the relevance of all claims to change, and that we need to start our evaluation over from scratch because we have a new worst-off claimant. Whilst I am generally in agreement with this point, I just do not see how it justifies violating the *Principle of Net Addition* and *Greater Consideration for Stronger Claims*. In the *First Horn for Global Relevance* case, for instance, there is a new worst-off claimant, and so we might think we need to start our evaluation over to account for the special concern we have for the worst-off claimant. However, doing so on Steuwer’s view,

Match in Order of Strength

undermines the worst-off claimant, such that the view fails to achieve its stated goal. Even if we do not think this is fatal, a view that can avoid violating these principles will be much stronger.

Match in Order of Strength

So, we have seen that these views all go wrong, more or less in the same way; they all have implausible conclusions in certain cases. I believe this is sufficient to dismiss these cases, especially if I can develop an account that does not violate any of these principle (which I do in the next chapter). However, it would be useful to dive a little deeper to diagnose why exactly Van Gils and Tomlin's and Steuwer's views go wrong. This will both strengthen my arguments against these views and it will help us to understand what a successful principle must do instead.

To do so we need to understand the reasoning behind *Sequential Claims Matching* and *Hybrid Balance Relevant Claims* and what makes them attractive in the first place. Now, as we saw above, both views share the same core idea for the order in which we ought to match claims. I believe this idea can be summed up by the following matching principle:

Match in Order of Strength: Claims should be matched in order of strength. I.e., we match the strongest claim first and then the second strongest claim, and so on.

Van Gils and Tomlin introduce the idea as a way to give priority to the worst-off claimant. As we have seen above, they argue that “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” (Van Gils & Tomlin, 2019, p. 228).

Steuwer expands on the idea:

Not saving a person with a strong claim to be saved requires a special justification to this person. Given that the person with the strongest claim is most likely to have grounds for

grievance or complaint against our failure to save her, our justification for acting as we do must be primarily addressed to those who have the strongest claim to our aid. This gives an answer to the question of the order of balancing. It explains why balancing starts with the strongest claim and then works its way down to less strong claims. The stronger the claim the greater is the urgency to give a justification to this person (Steuwer, 2021a, p. 24).

So, in essence, *Match in Order of Strength* is justified by a pro tanto reason to save the strongest claim. Now, there are two issues here. The first is whether there is a pro tanto reason to save the person with the strongest claim and that this claimant requires a special justification if they are not saved. The second is whether this pro tanto reason entails *Match in Order of Strength*.

I agree with my opponents on the first of these issues, it seems right to say that we have a pro tanto reason to save the person with the strongest claim and if that claim is not capable of being matched or outweighed, then we must save that claimant. This is the attractive element of the view, and what makes *Match in Order of Strength* seem plausible. This aspect of the view should be salvaged and should play a role in whichever matching principle we do decide on.

The second issue is where Steuwer's argument is flawed. We can see this in the last step of Steuwer's argument. He says, "The stronger the claim the greater is the *urgency* [emphasis added] to give a justification to this person" (Steuwer, 2021a, p. 24). But this is not quite right, a more accurate statement would be to say that the stronger the claim the greater the *importance* to give a justification to this person. But, of course, importance does not entail urgency, nor does it require that the strongest claim be matched first. If by matching other claims first we can better respect the importance of the strongest claim, then we ought to do so.

On Van Gils and Tomlin's, and Steuwer's views, when they check to see if the strongest claim can be matched, they do one of two things: (a) If the strongest claim cannot be matched, then they end the deliberation and save the group with the strongest claim, or (b) if the claim can be matched, then that claim is matched and set aside. Now, (a) is the correct move: (a) captures the pro tanto reason

Match in Order of Strength

to save the strongest claimant - if the strongest claim cannot be matched then we must save that claimant. However, (b) is not necessary to capture the pro tanto reason to save the strongest claimant, and whilst it seems the obvious move, there is no reason to assume that this is the best move - there may be more appropriate matchings.

Furthermore, if the pro tanto reason to save the strongest claimant is fleshed out in this way, it is in fact detrimental to that claimant. To see this let us revisit the First Horn case again:

First Horn

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia		
3 - Lost Limb	(100)	(6)
4 - Lost Hand		1000

At Stage 1 the strongest claimant is saved because the thousand B4 claims are not relevant. But at Stage 2, when we add equal strength claims to both groups but more to Group A, we swap away from saving the strongest claimant. This happens not despite of, but because of, *Match in the Order of Strength*: the A1 claim matches with the B3 claims, leaving the remaining B4 claims to outweigh the A3 claims. By prioritising matching the strongest claim, we in fact dismiss the strongest claim at the earliest opportunity. By urgently trying to give the strongest claimant a justification, we remove the strongest claimant from consideration at the earliest opportunity. Thus, we give the strongest claimant less of an opportunity to influence whom we ought to save. In this way, *Match in Order of Strength* rather than respecting the importance of the strongest claim seems to undermine its importance!

A matching principle that could incorporate (a) (i.e., checking whether the strongest claim can be matched) without (b), and promote the strongest claim in cases such as the above, would better respect this pro tanto reason to save the strongest claimant. I will develop such a view in the next chapter.

Match in Reverse Order of Strength

Before we conclude this chapter, let us consider one potentially obvious response to *Match in Order of Strength*. If *Match in Order of Strength* dismisses the strongest claims from consideration at the earliest opportunity, undermining said claimants' abilities to affect the outcome, then we might think that a view which keeps the strongest claims in competition for as long as possible would be a good solution. Such a view, we might think, would give the stronger claims the greatest chance of being the decisive claims.

The following principle would capture this idea:

Match in Reverse Order of Strength: Claims should be matched in reverse order of strength.

I.e., we match the weakest claim first and then the second weakest claim, and so on.

Of course, this principle only tells us the order in which we ought to match claims and not which claims we ought to match claims to. So, much like *Match in Order of Strength*, it would need to be supplemented by some further matching principles.

However, before we even get to that stage, we can see that *Match in Reverse Order of Strength* lands on the First Horn of Horton's dilemma and thus violates both the *Principle of Net Addition* and *Greater Consideration for Stronger Claims*. For the sake of brevity, *First Horn Sharpened* should be sufficient to demonstrate this:

First Horn Sharpened

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia	(100)	
3 - Lost Limb		(6)
4 - Lost Hand		1000

In this case, at Stage 1 we save Group A. At Stage 2, with the addition of the new claims, we first match the lost hand claims with the paraplegia claims. This leaves the lost limb claims to outweigh the death claim, and so we swap to saving Group B. Thus, leaving the strongest claimant in the

Concluding Remarks

competition for as long as possible will not necessarily be sufficient to favour the stronger claims.

We will still end up swapping away from the stronger claims in these cases.

Concluding Remarks

So, where does this leave us? I have shown that the matching principles provided in the literature have significant flaws, that they ought to be abandoned and that Local Relevance will need to consider different matching principles. This might not seem like much positive progress, but we have also identified the criteria a more successful Local Relevance view will need to fulfil. This will help us to develop a stronger Local Relevance view in the next chapter.

Such a view should respect *Equal Consideration for Equal Claims*, *Greater Consideration for Stronger Claims*, and the *Principle of Net Addition*. It should also respect the strongest claimant's pro tanto claim to be saved in such a way that is not detrimental to the strongest claimant and it should capture the special importance of giving the strongest claimant a justification when she is not saved. It should answer both the order in which to match claims and which claims should be matched with which, such that there are no ambiguities. Lastly, of course, there should be good independent reason to sustain it.

At minimum, this chapter shows that more attention is needed on the issue of matching principles and their importance to limited aggregation views. Even if one disagrees with my conclusions, I hope this chapter goes some way to highlighting that importance.

Chapter 3 – Match by Closeness

Introduction

In the last chapter, I showed the importance for Local Relevance views of getting the right matching principle and that none of the options so far offered do so. In doing so we identified where such matching principles went wrong and that a more successful matching principle will need to avoid a number of problems. It will need to respect *Equal Consideration for Equal Claims*, *Greater Consideration for Stronger Claims*, and the *Principle of Net Addition*. It should also respect the strongest claimant's pro tanto claim to be saved and the special importance of giving the strongest claimant a justification when she is not saved. However, it must do so in a way that is not detrimental to the strongest claimant's chances of being saved. We also saw that the view needs to answer both the order we should match claims in, and which claims we should match to which if it is to avoid ambiguous cases.

Of course, beyond simply avoiding the problems identified in the previous chapter there should be positive reasons to adopt any alternative matching principle. Thus, there needs to be a good independent justification behind such a matching principle, such that the view is not simply an ad hoc solution. Lastly, we should also expect any matching principle to navigate an escape from Joe Horton's dilemma in a convincing manner. In this chapter I will develop a principle that I believe can meet all these criteria.

Matching by Closeness

The answer I propose is to consider matching similar claims first. I shall call this matching principle

Match by Closeness:

Match by Closeness: when matching claims, where possible match claims with closest relevant claims.

By 'closeness' I mean 'closeness of strength', a claim against a lost arm is closer in strength to a claim against a lost hand than it is to a claim against death. Thus, if we have a choice whether to match lost arm claims with lost hand or death claims, *Match by Closeness* tells us to match the lost arm claims with the lost hand claims. Already, this seems highly intuitive.

Importantly, this principle tells us both (a) the order in which we ought to match claims and (b) which claims we should match each particular claim to. Firstly, we match claims in order of closeness. So, we identify the claim(s) with the closest relevant competing claims and match these claims first. We then identify the remaining claims with the next closest relevant competing claims and match these claims, and so on. Secondly, of course, we also match the identified claims to the closest competing claims still under consideration. Thus, there should be no ambiguity as to how to match the claims.²²

Now this view might sound quite complicated. Thus, it might be helpful to see how the view works when it is broken down. To do so I think it is helpful to borrow the sequential element of Van Gils and Tomlin's *Sequential Claims Matching* view (Van Gils & Tomlin, 2019, pp. 228-231) and amend it such that it matches by closeness rather than in order of strength. This sequence will match first identical claims, then similar claims, and then wider relevant claims, until we reach the limit of relevance. The following is such a sequence (also illustrated as a flow chart):

²² An ambiguity does actually arise when there are equidistant competing claims. We will return to this ambiguity later in this chapter.

*Sequential Claims-Matching by Closeness*²³

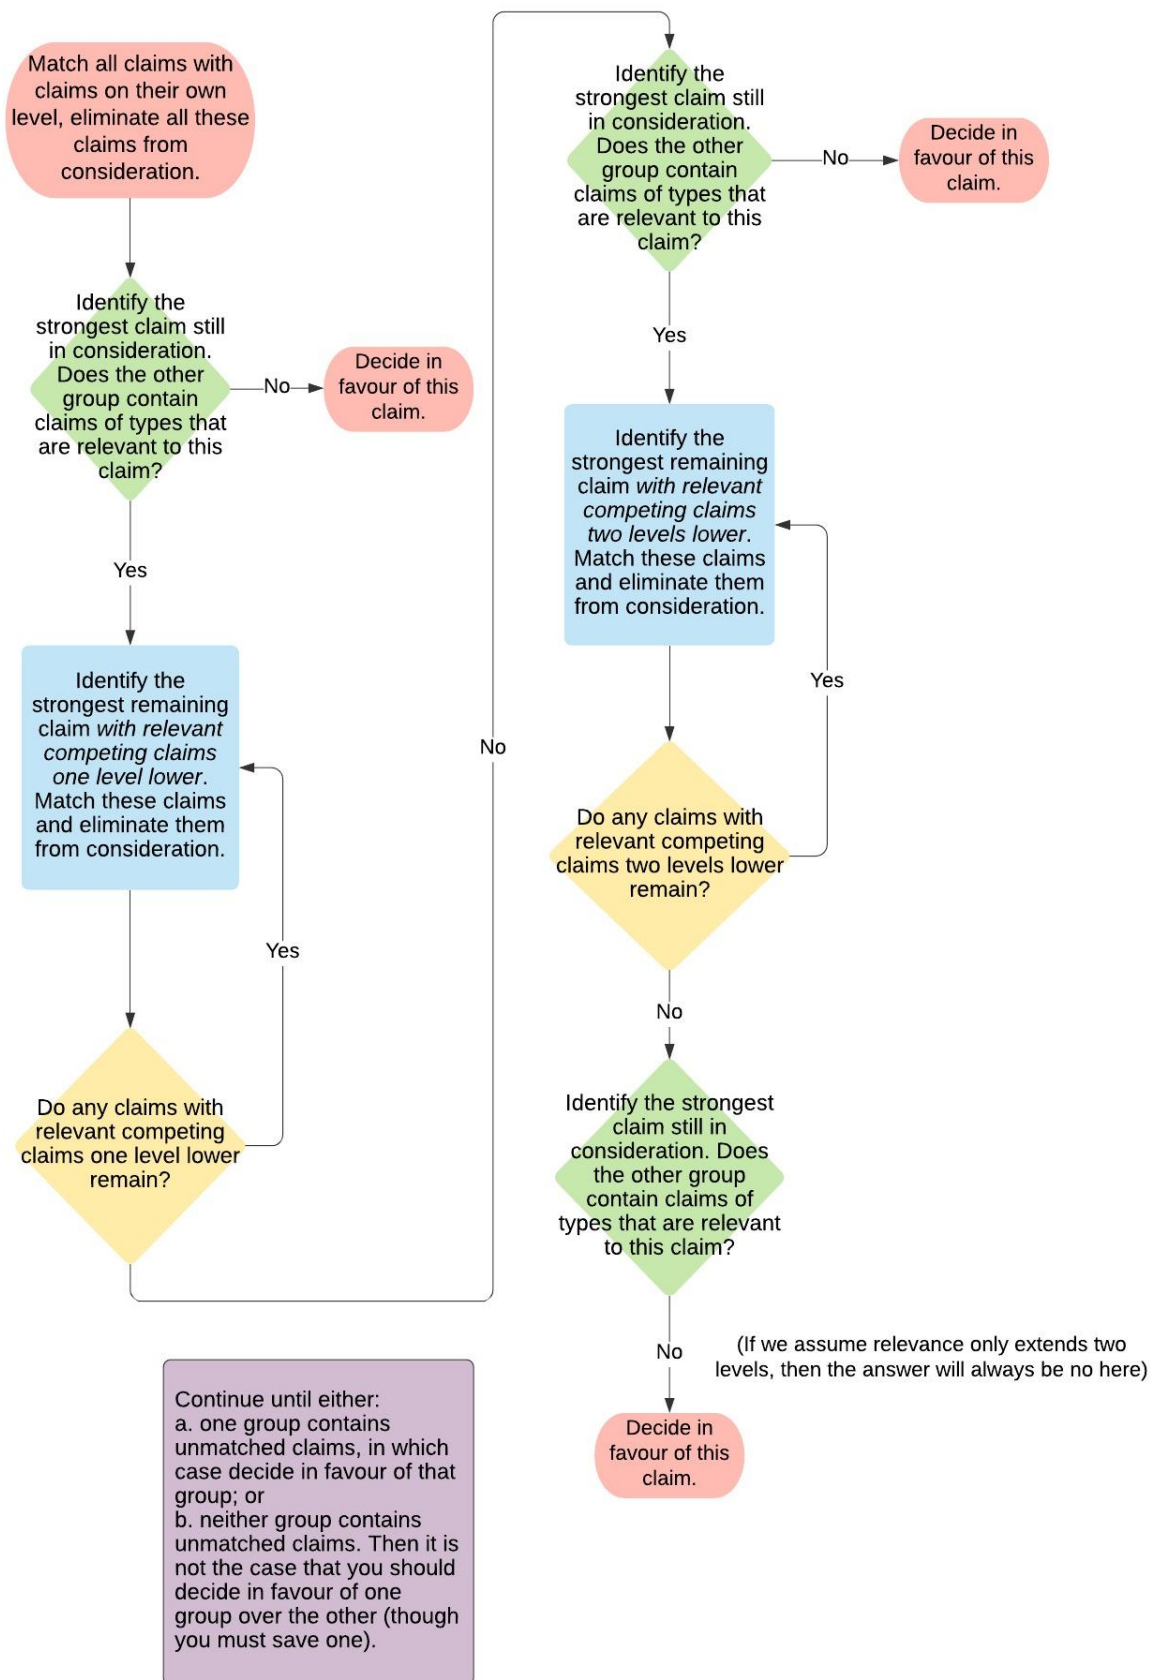
- I. Match all claims with claims on their own level, eliminate all these claims from consideration.
- II. Identify the strongest claim still in consideration. Does the other group contain claims of types that are relevant (i.e., sufficiently strong relative) to this claim and sufficient in number to match it?
 - a. If not, decide in favour of the group with the strongest claim still in consideration.
 - b. If so, go to Step III.
- III. Identify the strongest remaining claim *with relevant competing claims one level lower*. Match these claims [with the claims one level lower] and eliminate them from consideration.
- IV. Repeat Step III until there are no claims with relevant competing claims one level lower.
- V. Identify the strongest claim still in consideration. Does the other group contain claims of types that are relevant to this claim and sufficient in number to match it?
 - a. If not, decide in favour of the group with the strongest claim still in consideration.
 - b. If so, go to Step VI.
- VI. Identify the strongest remaining claim *with relevant competing claims two levels lower*. Match these claims [with the claims two levels lower] and eliminate them from consideration.
- VII. Repeat Step VI until there are no claims with relevant competing claims two levels lower.
- VIII. Identify the strongest claim still in consideration. Does the other group contain claims of types that are relevant to this claim and sufficient in number to match it?²⁴
 - a. If not, decide in favour of the group with the strongest claim still in consideration.
- IX. Continue until either:
 - a. one group contains unmatched claims, in which case 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).

²³ A shorter version of this sequence can be found in Appendix 2. I use the longer version here both to match Van Gils and Tomlin's *Sequential Claims-Matching* formulation, and also for the sake of clarity when explaining the idea. I will use the shorter version in future chapters as it is less wordy and less complicated once we have a hang of the basic idea.

²⁴ At this point there will be no relevant claims left for any ranking where relevance only extends down two levels.

Matching by Closeness

Sequential Claims Matching by Closeness - Flow Chart



Chapter 3 – Match by Closeness

Now to see how *Sequential Claims Matching by Closeness* captures *Match by Closeness*, and how it handles particular cases, let us reconsider the test case from the previous chapter:

Case 8

Level	Group A	Group B
1 - Death	2	1
2 - Paraplegia		
3 - Lost Limb		5
4 - Lost Hand	4	

In this case, Group A contains two death claims and four lost hand claims. Group B contains one death claim and five lost limb claims. Remember from the previous chapters, that for illustrative purposes we have split claims into different levels and set relevance at two levels, such that the lost limb claims are relevant to death claims, but lost hand claims are not. We also decided that for ease of reference we would call the Level 1 claims in Group A “A1 claims”, Level 2 claims in Group B “B3 claims” and so on. Lastly, we agreed that two claims match with exactly one claim on the level above - so two paraplegia claims match exactly with one death claim. With this in view let us walk through *Sequential Claims Matching by Closeness* step by step:

Step I: Firstly, we identify all claims that have competing claims on the same level - i.e., identical competing claims. In this case, these are our death claims. We then match these claims until one group has no claims of this level. I have demonstrated this by striking through those matched claims and putting the remainder in square brackets, as per the table below:

Case 8

Level	Group A	Group B
1 - Death	2 [1]	1
2 - Paraplegia		
3 - Lost Limb		5
4 - Lost Hand	4	

Matching by Closeness

Step II: Next, we identify the strongest claim remaining in the competition. This is the remaining death claim, A1, and we look to see whether there are any claims relevant to the A1 claim. The B3 claims are relevant and sufficiently strong in number to outweigh the A1 claim. As such, we move on to Step III. (Note here that we do not match the A1 claim with the B3 claims, but instead move on to Step III. This is because Step II tells us nothing about which claims to match, but simply whether there are any claims relevant to the strongest claim, and whether we should carry on matching claims.)

Step III: Now we must identify the strongest remaining claims with competing claims one level lower. In this case these are the B3 claims, as the B3 claims have competing claims one level lower - i.e., the A4 claims. We then match the B3 claims with the A4 claims, leaving three B3 claims remaining:

Case 8

Level	Group A	Group B
1	2 [1]	1
2		
3		5 [3]
4	-4	

Step IV: This step tells us to repeat Step III until there are no more claims with competing claims one level lower. However, in this case there already are no more claims with competing claims one level lower and so we move straight to Step V.

Step V: This step repeats Step II. We identify the strongest remaining claim, the A1 claim, and check if there are relevant claims that could outweigh it. The three B3 claims are relevant but, given that four B3 claims are needed to match the A1 claim, we can see that the A1 claim cannot be matched. Thus, Step V tells us to save the group with the strongest claim still in consideration - i.e., Group A.

Hopefully, this should make clearer how *Match by Closeness* operates, and what makes the view distinctive from other views. Whilst it gets the same answer, in this case, as a view operating on

Match in Order of Strength (such as *Sequential Claims-Matching*) the way in which it does so is clearly different, as the strongest claim remains in the competition until the end.

Independent Justification

With the principle now in view, we can move on and analyse this principle in detail. Let us first consider the justification or independent reason we might have for this principle. Firstly, it seems to me that we have strong reason to treat claimants with equal claims equally. Now, of course, this is not always possible, and as Tomlin rightly says “any plausible view concerning whom we should save between competing groups will allow that people holding equal claims may be *treated differentially*. Unless you are prepared to save nobody, in any case with equal claims on both sides you will save some but not others with the same strength of claim.” (Tomlin, *On Limited Aggregation*, 2017, p. 242)

But it seems clear to me that in these cases it is still regrettable that we could not treat the claimants equally, and not just because it is regrettable that we could not save everyone. As Broome notes, there is a concern of fairness here (Broome, 1998).²⁵ What this indicates is that there is a pro tanto reason to treat claimants with equal claims equally. Now, this pro tanto reason can, of course, be outweighed by other reasons, such as when we can only save one of two people with equal claims. Nevertheless, much like the pro tanto reason to save the strongest claimant, it still maintains some force.

Match by Closeness captures this pro tanto reason. By prioritising matching equal claims, and then similar claims, we ensure that all equal claims are first matched with each other. This is captured in Step I of *Sequential Claims Matching by Closeness*, where we match all equal strength claims with each other. Only if there are remaining claims after this initial matching do we treat any of the equal

²⁵ We will return to Broome’s view and questions about fairness and equal treatment in Chapter 5. This very brief treatment should be sufficient for now to see the motivation behind *Match by Closeness*.

Independent Justification

claims differently by matching them with other claims. This seems like a good way to respect the pro tanto reason to treat claimants with equal claims equally, even when we cannot save all the equal claimants. In this way, *Match by Closeness* necessarily avoids violating *Equal Consideration for Equal Claims*.

Furthermore, we can see that *Match by Closeness* seems like a natural extension of what Kamm calls the *Substitution of Equivalent*s (Kamm, 1998, Chapter 6). *Substitution of Equivalent*s recognises that “it is morally permissible to balance off equal and opposing individual claims or needs [... and] that neither of two equal and opposing claims or needs can finally decide an outcome, the “unbalanced” members of one side must do that” (Kamm, 1998, p. 101). To see why, let us reconsider a familiar case:

First Horn

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia		
3 - Lost Limb	(100)	(6)
4 - Lost Hand		1000

In this case we have added equal strength claims to both sides but in differential numbers. Now if we change to saving Group B at Stage 2, then the B3 claims have decided the outcome. But *Substitution of Equivalent*s indicates, intuitively, that no claim with an equal and opposing claim should be decisive.²⁶ Instead, Kamm suggests we should set aside equal strength claims until at least one group is entirely set aside, such that only the unbalanced claims can be decisive. Whilst we might be able to balance the death claim with lost limb claims, we would better respect the claimants if we balanced the lost limb claims with other lost limb claims. This, of course, is exactly what *Match by Closeness* suggests.

²⁶ In this way *Substitution of Equivalent*s is a mirror of *Equal Consideration for Equal Claims*.

Match by Closeness simply takes the logic of *Substitution of Equivalentents* a step further. It widens this approach to near-equivalentents. *Match by Closeness* suggests that it would be better to balance lost limb claims with lost hand claims, than with death claims, because the lost hand claims are closer in strength. Doing so is more respectful to the claimants. Thus, *Match by Closeness* says that we should substitute equivalentents where possible, and balance similar claims before comparing more distant claims.

Match by Closeness also tracks with other deeper justifications of relevance views. As previously mentioned, for reasons of space and to keep my conclusions general, I do not want to commit myself to any particular justification here. Nevertheless, I want to quickly outline how *Match by Closeness* fits with the common thread running through these different justifications. Voorhoeve (2014), Kamm (1998), Scanlon (1998), and Nagel (1995) all approach the justification of relevance views as an attempt to balance partial and impartial concerns or the personal and impersonal perspectives. To put it very simply, we must take up the personal perspective of each individual claimant to see what is at stake for them, and what is permissible and respectful from each of their perspectives. For our purposes, I think the simplest way to express this is through Steuwer's Respectful Failure to Save Principle:

Respectful Failure to Save Principle. Every person whom we fail to save is entitled to a respectful justification for our failure to save. It is disrespectful and impermissible to fail to save a person with a strong claim for the sake of persons whose claims are irrelevant to this strong claim. (Steuwer, 2021a, p. 24)

In other words, a claim is relevant if it could form part of a respectful justification (from the perspective of the competing claim) not to save a competing claim. With this in view we can see how *Match by Closeness* can adhere to the personal perspective.

Match by Closeness does a better job of both (a) giving each matched claimant a better explanation as to why they have been taken out of consideration and (b) giving each losing claimant a strong and

Independent Justification

respectful justification as to why they have not been saved. This is because the justification given to each claimant is more likely to refer to other similar claims. Thus, a paraplegia claimant would be given a justification referencing other paraplegics or similarly strong claims, rather than weaker claims such as lost hands. Such a justification will be more acceptable from the claimant's perspective and thus more respectful of her claim. At its most basic it is more respectful to say to someone 'I am sorry we cannot save you from paraplegia, because we are busy saving others from paraplegia' than it is to say 'I am sorry we cannot save you from paraplegia, because we are busy saving others from lost hands'.

What this also indicates is that such a matching principle can capture how our respect-based reasons are scalar rather than binary. On Van Gils and Tomlin's, and Steuwer's, views, it is either respectful or disrespectful to compare a claim. If it is respectful to compare a claim, it is relevant; and if it is disrespectful to compare a claim, it is irrelevant. Respect plays no role in deciding how to match claims beyond deciding whether a claim is relevant or not.

But *Match by Closeness* allows us to capture these respect-based reasons, even between relevant claims: it is a little disrespectful to compare one person's death claim to other people's paraplegia claims, but it is more disrespectful to compare a death claim to lost limb claims (though in neither case disrespectful enough to rule out such comparisons).²⁷ Thus, it is best to compare a death claim to paraplegia claims rather than lost limb claims. Relevance is merely the success condition: at some point the comparison becomes so disrespectful that the weaker claims are completely silenced and so we count those claims as irrelevant.

Match by Closeness can respect the personal perspective in another way too. Consider the First Horn case again:

²⁷ Even comparing claims of the exact same strength seems a little disrespectful given the separateness of persons and what is at stake for each of the claimants.

Chapter 3 – Match by Closeness

First Horn

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia		
3 - Lost Limb	(100)	(6)
4 - Lost Hand		1000

On views like Van Gils and Tomlin's, and Steuwer's, lost limb claims in Group B are being matched with death claims, whilst lost limb claims in Group A are being matched with lost hand claims. Claims of the same strength are being compared to two set of claims of very different strength. From the personal perspective of the lost limb claimants in Group A this matching will feel very unfair. They would have good reason to complain that their claims are not being taken as seriously as the lost limb claimants in Group B. *Match by Closeness* prevents this from happening and ensures the lost limb claims are matched with each other. This will be acceptable from the personal perspective of both groups, with neither group able to complain that their claims are not being treated fairly.

Lastly, although the focus of *Match by Closeness* is as a way to capture a theory of right, I want to quickly touch on why we might also have good decision-theoretic reasons for *Match by Closeness* too. When making decisions it is easier to compare identical or similar claims than it is to compare dissimilar claims. It is very easy to see that a claim against paraplegia is as strong as another person's claim against paraplegia. It is less easy to compare the varying strengths of claims against quadriplegia and against paraplegia - how much worse is quadriplegia, is the quadriplegia claim 1.5x stronger, 2x stronger? It is harder still to compare across even more dissimilar claims - is death 4x, 10x, 100x worse than paraplegia? For this reason, when we are comparing groups of claims, it makes good decision theoretic sense to prioritise matching by closeness - this way we are less prone to error. Whilst this tells us nothing about the veracity of such a procedure, it is nevertheless advantageous that this procedure limits the number of difficult comparisons.

So, there is ample rationale behind *Match by Closeness*, sufficient at least for our purposes. More could be said to tie it to a particular account of relevance, but as I do not want to commit myself to

Testing Match by Closeness

any particular view here, I shall leave that discussion for others. *Match by Closeness* is a broad enough principle that I see no reason that it would not be amenable to any of the major limited aggregation rationales. Hopefully, by presenting it in general terms, I make apparent the generality of my conclusions. If it turns out that *Match by Closeness* is incompatible with a particular account of the deeper justification of limited aggregation views, then we can return to the subject. Until then, I think we have sufficiently established *Match by Closeness* and are safe to proceed.

Testing Match by Closeness

With the justification in view now let us turn to how *Match by Closeness* avoids the problems of the previous views. We will consider how it captures the pro tanto reason to save the strongest claim, without being detrimental to the strongest claim, and how it avoids violating both *Greater Consideration for Stronger Claims* and the *Principle of Net Addition*.

Respecting the Strongest Claim

Firstly, let us look at how *Match by Closeness* captures the pro tanto reason to save the strongest claim. *Match by Closeness* maintains that if the strongest claim in the competition cannot be matched by any other claims, then we automatically save the group with the strongest claim. By fleshing this pro tanto reason out this way we also avoid it being detrimental to the strongest claimant. Consider again *First Horn Sharpened*:

First Horn Sharpened

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia	(100)	
3 - Lost Limb		(6)
4 - Lost Hand		1000

At Stage 1, of course, Group A wins. At Stage 2, we add more and stronger claims to Group A. Now recall from the previous chapter that Van Gils and Tomlin's view is supposedly derived from this pro tanto reason to save the strongest claim, yet the view they develop ends up swapping to save Group

B at Stage 2 and is thus detrimental to the strongest claim. This is because their view involves two steps: (a) If the strongest claim cannot be matched, then they end the deliberation and save the group with the strongest claim, or (b) if the claim can be matched, then that claim is matched and set aside. But we saw that only (a) is necessary to capture the pro tanto reason to save the strongest claim, whilst (b) undermines it in cases like *First Horn Sharpened*; (b) tells us to match the A1 claim with the B3 claims, which then allow the B4 claims to outweigh the A2 claim.

Match by Closeness does not make such a mistake; it maintains (a), respecting the pro tanto reason to save the strongest claim, but rejects (b). It is easiest to see why by going through *Sequential Claims Matching by Closeness*.

Step I: there are no claims on the same level so we skip this step.

Step II: we identify the strongest claim, A1, and can see that there are relevant B3 claims strong enough to outweigh the A1 claim. Therefore, we move to Step III.

Step III: we identify the strongest claims with claims one level lower. These are the A2 claims, and we match these claims with the closest relevant claims, the B3 claims, and set these aside. This is indicated on the table below:

First Horn Sharpened

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia	(100) [97]	
3 - Lost Limb		(6)
4 - Lost Hand		1000

Step IV: this step tells us to repeat Step III with the next strongest claim with competing claims one level lower - but as there are none, we move on to Step V.

Step V: we identify the strongest remaining claim, A1, and check to see if there are relevant competing claims that can outweigh it. There are not, and so *Sequential Claims Matching by Closeness* tells us to save Group A.

Testing Match by Closeness

Steps II and V are the steps here that capture (a) and the pro tanto reason to save the strongest claim. These steps tell us to check whether there are relevant competing claims to the strongest claim. If there are none, then these steps tell us to save the strongest claimant. Thus, we capture the special importance of the strongest claim and the importance of saving them if their claim cannot be outweighed. However, by telling us to move on to Step III, rather than match the A1 claim immediately, we also avoid taking the strongest claim out of competition early, and so avoid matching in such a way that is detrimental to the strongest claim.

First Horn

Nevertheless, avoiding these consequences in these cases is not enough, as we might come up with other cases in which the underlying principles are violated. Thus, we need to see that *Match by Closeness* avoiding the problematic consequences of the First Horn, and the sharpened version of the First Horn, is no coincidence or artifact of these particular examples either. We need to see that *Match by Closeness* necessarily prevents the breaking of the *Principle of Net Addition* and *Greater Consideration for Stronger Claims*.

The *Principle of Net Addition* is not violated because when two groups of claims with the same strength but different numbers are added to any competition, the smaller group of claims is fully matched with the larger group. Thus, the smaller group is eliminated from consideration before it can affect the rest of the competition.

Similarly, *Match by Closeness* also always avoids violating *Greater Consideration for Stronger Claims*. This is because, when two groups of claims with the same numbers but different strength (or greater numbers for the stronger claims) are added to any competition, the weaker claims will be fully matched against the stronger claims. Either that, or the weaker claims will be fully matched against some other weaker claims leaving the stronger claims still in the competition. Thus, the weaker claims are taken out of consideration before they can influence the overall decision.

In the First Horn cases the stronger group is strengthened by either exactly equal or dominant claims at Stage 2, and this is why we think we ought not to swap whom we save.²⁸ *Match by Closeness* however respects the *Principle of Added Dominance*:

Principle of Added Dominance: adding claims cannot make the group to which dominant claims are added less choice-worthy compared with the group to which the dominated claims are added.

Views that violate this principle do so because they allow the dominated claims to affect the relevance of other claims. For instance, in *First Horn Sharpened* the dominated lost limb claims make the lost hand claims relevant by taking the death claim out of consideration.

With this in view, we can see how *Match by Closeness* necessarily avoids violating the *Principle of Added Dominance*. It does so because the dominated claims are by necessity fully matched either by the dominant claims themselves, or some weaker claims, such that they cannot match with the strongest claims and alter the relevance of other claims in the competition. In other words, *Match by Closeness* prevents dominated claims from affecting the relevance of other claims. Thus, *Match by Closeness* does not just get the right answers in these particular cases but will always avoid landing on the First Horn.

Second Horn

There is one last issue I ought to look at here. Whilst I argued in Chapter 1 that the First Horn of Horton's dilemma is fatal, I also argued that the second horn, whilst unavoidable, is not fatal if it is handled in the right way. So, it would be helpful to see how *Match by Closeness* handles the Second Horn. To demonstrate this horn, consider again the Second Horn:

²⁸ Return to Chapter 1 for more detail on this.

Testing Match by Closeness

Second Horn

Level	Group A + C	Group B + D
1 - Death	1	
2 - Paraplegia		
3 - Lost Limb	100	6
4 - Lost Hand		1000

As we have already seen, *Match by Closeness* tells us to save Group A in this case and thus avoids the First Horn. However, the Second Horn can be split into two separate pairwise comparisons,

Cases 1 and 2:

Case 1

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia		
3 - Lost Limb		6
4 - Lost Hand		

Case 2

Level	Group C	Group D
1 - Death		
2 - Paraplegia		
3 - Lost Limb	100	
4 - Lost Hand		1000

In Case 1 *Match by Closeness* tells us to save Group B. In Case 2 *Match by Closeness* tells us to save Group D. Yet in the Second Horn case, when the groups are combined, *Match by Closeness* tells us to save Group A+C. By saving Group A+C without changing the strength or number of claims, we turn 'loser-groups' into 'winner-groups' and 'winner-groups' into 'loser-groups'. This is a slightly bizarre implication: it does not seem that combining claims in this way should change who wins.

Yet, whilst it does seem like a bizarre implication, there is good reason to think this horn is not devastating. Firstly, as we saw in the first chapter, the Second Horn of the dilemma is unavoidable for all relevance views. This is because there are some cases which, depending on the way they are

divided, will make either group the ‘winner-groups’. More importantly we saw that the Second Horn is premised on a faulty assumption of *Weak Additivity*:

Weak Additivity: If A is preferable to B, and C is preferable to D, then A with C is preferable to B with D.

Weak Additivity fails whenever there are interaction effects between elements, such as in cases of substitute or complimentary goods, and that is exactly what Local Relevance theorists say occur in these cases. The presence of a stronger claim might interact with weaker claims such as to make the weaker claims irrelevant.

Furthermore, when Cases 1 and 2 are treated separately the decision we make in one case does not restrict who we can save in the other case. However, when they are combined into the Second Horn Case, we are restricted. For instance, we could choose to save the lost hand claimants in Case 2 and still choose to save the death claimant in Case 1, but when the cases are combined saving the lost hand claimants will restrict our choice in Case 1, such that we can no longer save the death claimant. Thus, we should not think that the Second Horn is devastating to Local Relevance views.

Nonetheless, we still need an explanation for why combining the two ‘loser-groups’ can turn them into the winning group. What moral features mean that, when Cases 1 and 2 are combined, we ought to start our evaluation over? It is here that *Match by Closeness* helps to explain why we would want to swap to the two ‘loser’ groups in the Second Horn Case. The qualitative shape and what we can influence has changed, meaning what is most respectful has also changed.

In Case 1 the B3 claims had to be matched against the A1 claim because there was no other choice. Of course, this involves comparing a death claim with lost limb claims, which is not ideal, as these claims are quite different in strength. When Cases 1 and 2 are combined into the Second Horn Case, however, a new potential matching arises. We could match the B3 claims with the A3 claims, and compare lost limb claims with lost limb claims. This, of course, would be a much more respectful

Ambiguity

comparison of claims. If we maintained ties to the first way of comparing and matching claims, then the lost hand claimants (especially those in Group A+C) would have a strong claim to being treated unequally and unfairly.

Thus, *Match by Closeness* can give us independently good reason to swap how we match claims in the Second Horn Cases. The swap is justified because a more appropriate and respectful way to match claims arises.

Ambiguity

Now, whilst *Match by Closeness* does answer both the order in which we should match claims and which claims each particular claim should be matched to, an ambiguity still remains in certain cases. Namely *Match by Closeness* cannot tell us how to match a claim when that claim is equally close to two competing sets of claims. To see why consider Case 10:

Case 10

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia		2
3 - Lost Limb	4	
4 - Lost Hand		10

In this case the B2 claims are equidistant between the A1 claims and the A3 claims. The A3 claims are also equidistant between the B2 and B4 claims. So, there are two ambiguities here: firstly, which claims ought we to match first, and secondly, which claims ought we to match each particular claim to.

If the A1 claims are matched first with the B2 claim, then Group B wins. On the other hand, if the B2 claims are matched first, and are matched with the A3 claims, then Group A wins. Unfortunately, *Match by Closeness* has no answer as to how we match these claims.

Of course, any relevance view will need an answer in this case and so *Match by Closeness* will need to be supplemented by some other matching principle. I think there are a number of ways to settle this ambiguity. The following are four possible supplementary principles. The first three are derived from principles considered in the previous chapter. The last is a slightly different approach to ambiguous cases offered by Van Gils and Tomlin:

Strongest Decides: when there is an ambiguity over how to match up claims, then the ambiguity should be settled in whatever way is in the interest of strongest overall claimant.

Match in Order of Strength: when there is an ambiguity over how to match up claims, then we should settle the ambiguity by matching the ambiguous claims in order of strength. I.e., we match the strongest ambiguous claim first and then the second strongest claim, and so on.

Match in Reverse Order of Strength: when there is an ambiguity over how to match up claims, then we should settle the ambiguity by matching the ambiguous claims in reverse order of strength. I.e., we match the weakest ambiguous claim first and then the second weakest claim, and so on.

“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).” (Van Gils & Tomlin, 2019, pp. 246-247)

Because these principles are only telling us how to match when *Match by Closeness* has no answer, they do not land on the First Horn of Horton’s dilemma. Thus, we cannot use *Equal Consideration for Equal Claims*, *Greater Consideration for Stronger Claims* nor the *Principle of Net Addition* to decide which supplementary matching principle to use. We can only decide which matching principle to use based on the thought process behind each one, and what we want to give priority to.

Strongest Decides

Let us start with *Strongest Decides*. *Strongest Decides* in Case 10 will tell us to match the B2 claims with the A3 claims, as doing so will mean that Group A and the strongest claimant will win. This, of course, gives priority to the strongest claimant, which fits well with the limited aggregation approach. It treats the strongest claimant as if she has the deciding vote in such cases.

However, there are a number of drawbacks to *Strongest Decides*. Firstly, we might think that allowing the strongest claimant the power to decide how to match claims is to give her too much power. It might seem like the strongest claimant can gerrymander the way we aggregate the claims such that her claim can win. More importantly, this might indicate that the other claims are not being treated with the due respect which they deserve. Weaker claims are being matched with each other based on what is favourable to another claim, rather than what is most respectful to each of these claimants. Thus, it seems like the strongest claimant might be treating them as means to an end. Furthermore, this view seems like an ad hoc and slightly arbitrary way to decide. It will not be clear from the outset how we will be matching claims and so the process might not seem like one all the claimants could agree on. We might think it would be preferable to have a principle set out at the beginning which can be agreed to by all claimants, and which settles these ambiguities in advance.

Secondly, *Strongest Decides* will not completely solve the ambiguity anyway. Consider the following case:

Case 11

Level	Group A	Group B
1 - Death	2	1
2 - Paraplegia		2
3 - Lost Limb	4	
4 - Lost Hand		10

In Case 11, there is no overall strongest claimant, and so we are still left with an ambiguity. The response to this might be to say that we should match in the interest of the strongest claimant still

under consideration. Given that two of the death claims in this case match, the remaining death claim in Group A should get the deciding vote. However, this seems to move us further away from the initially intuitive idea behind *Strongest Decides*.

Furthermore, there might be cases in which there are equal numbers of the strongest claimants, and the ambiguity arises again. If we were to then say that we ought to let the next strongest overall claimant decide then we seem to move even further away from the intuitive idea behind *Strongest Decides*. Such cases might also worsen the feeling that the deciding claim is being given too much power here, and that this way of settling the ambiguity is too ad hoc.

Thirdly, if we match in favour of the strongest claimant in Case 10, then we match the B2 claims with the A3 claims. However, we might think this to be less than optimal in this case. Doing so would mean that the lost hand claims must be compared with the death claims. We are forced into comparing distant claims when it is preferable to compare similar claims.

If instead we compared the B2 claims with the A1 claim, then we could also match very similar claims with A3 and B4. Thus, it might seem that *Strongest Decides* goes against *Match by Closeness* in the sense that it reduces the opportunities to match claims to their closest competing claims. I think this is the strongest reason to reject *Strongest Decides*. It seems to go against the ethos of *Match by Closeness*.

Lastly, because *Strongest Decides* matches claims in this way it will also increase the number of claims that will be considered irrelevant. In fact, this is how *Strongest Decides* will ensure the strongest claimant will be saved, by matching in such a way that the weakest claims are considered irrelevant. We might also think that this is a mark against *Strongest Decides*. Whilst it is, of course, expected that claims might become irrelevant when matching, it seems like we should not operate with the intention of making claims irrelevant. But this is exactly what *Strongest Decides* does.

Match in Order of Strength and Match in Reverse Order of Strength

Now let us consider *Match in Order of Strength* and *Match in Reverse Order of Strength* together.

Both of these supplementary matching principles tell us to save Group B in Case 10:

Case 10

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia		2
3 - Lost Limb	4	
4 - Lost Hand		10

Match in Order of Strength tells us to settle the ambiguity by matching the A1 claim first. We match the A1 claim to the B2 claims and so the A3 claims are then matched and outweighed by the B4 claims. Thus, *Match in Order of Strength* tells us to save Group B.

Match in Reverse Order of Strength tells us to settle the ambiguity by matching the B4 claims first. We match the B4 claims to the A3 claims and then we match the B2 claims to the A1 claim. This leaves two unmatched B4 claims remaining, and so *Match in Reverse Order of Strength* tells us to save Group B.

So, what are the advantages of these approaches? Firstly, both these principles avoid the issues that affect *Strongest Decides*. Both views decide in advance the order in which claims will be matched if there is an ambiguity. This avoids giving the strongest claimant too much say over whom we save and prevents the gerrymandering of cases. Such views are thus more justifiable to each claimant, as claimants would have no reason to object to either principle without knowing their position in each case. This also means that we remove the chance of claims being deliberately compared in ways that make them irrelevant.

Furthermore, both principles better respect the ethos behind *Match by Closeness*. For instance, in Case 8 both views tell us to match the A1 claim with the B2 claims, and the A3 claims with the B4 claims. Thus, all claims are being compared to fairly similar claims, rather than only some being

Chapter 3 – Match by Closeness

compared with similar claims, and others being compared with more distant claims. I think the above points are sufficient to note that either of these principles is preferable to *Strongest Decides*.

How do we choose between these two supplementary principles then? Well firstly, it is important to see the cases in which they will give different answers. Consider Cases 12 and 13:

Case 12

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia		2
3 - Lost Limb	4	
4 - Lost Hand		
5 - Lost Finger		100

Case 13

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia		
3 - Lost Limb		4
4 - Lost Hand	8	
5 - Lost Finger		100

In Case 12, *Match in Order of Strength* tells us to match the A1 claim to the B2 claims first, and thus Group B wins. *Match in Reverse Order of Strength* conversely, tells us to match the A3 claims to the B2 claims first, and thus Group A wins.

In Case 13, *Match in Order of Strength* tells us to match the B3 claims to the A4 claims first, and thus Group A wins. *Match in Reverse Order of Strength* conversely, tells us to match the B5 claims to the A4 claims first, and thus Group B wins.

So, in these cases the views give different answers. In Case 12, *Match in Reverse Order of Strength* favours the strongest claimant, whilst *Match in Order of Strength* takes the strongest claimant out of competition earlier. However, in Case 13 this is reversed; *Match in Order of Strength* favours the

Ambiguity

strongest claimant, whilst *Match in Reverse Order of Strength* allows the strongest claimant to be outweighed.

Importantly, this means we cannot decide between these two supplementary principles based on which view favours the strongest claimant. More generally, I think this indicates that we cannot decide between these two views based on how they handle particular cases. Both views seem to give acceptable judgements in these cases and neither view violates any deeper principles. Perhaps, then, either supplementary principle could be adopted, so long as the principle is agreed upon in advance. If we do want to decide between the two supplementary principles, then we will need to do so entirely on the basis of the reasoning for either principle.

Firstly, let us look at the reasons in favour of *Match in Order of Strength*. As we saw in the previous chapter Steuwer argues that *Match in Order of Strength* would give a greater urgency to matching the strongest claim (Steuwer, 2021a, p. 24). But we also saw that the pro tanto reason to save the strongest claim, or the special justification owed to the strongest claimant, does not support matching the strongest claim with added urgency as doing so would sometimes run counter to the strongest claimant's interests. So, I do not think this can be the reason we ought to start with the strongest ambiguous claims.

Nevertheless, *Match in Order of Strength* should not be ruled out for these reasons: we might have other reasons to use it in ambiguous cases. One such reason is that *Match in Order of Strength* might be a more respectful way to match claims when there is an ambiguity. To see why consider Case 12 again:

Case 12

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia		2
3 - Lost Limb	4	
4 - Lost Hand		
5 - Lost Finger		100

In this case the ambiguity arises because the paraplegia claims are equidistant between the death claim and the lost limb claims. So, in essence the question is about how we match the paraplegia claims. We might think it is more respectful to match the paraplegia claims with the death claims in this situation. This is because it gives a better justification to the paraplegia claimants as to why they are being taken out of consideration. In other words, if the paraplegia claimant were to choose whose claims to match theirs with, they would choose the stronger.

Secondly, we might think that addressing the stronger claims first in ambiguous cases is intuitive in and of itself. Rather than derive this rule from a pro tanto reason to save the strongest claim, or from the special justification owed to the strongest claimant in particular, we might just think that we ought to be primarily focused on the stronger claims in general. Note how this is different from Van Gils and Tomlin's, and Steuwer's, justifications. It is not focused on giving priority to the strongest overall claim in particular, but on prioritising our focus onto the stronger claims as a whole. Importantly, this justification can side-step the accusation that *Match in Order of Strength* is detrimental to the strongest claim - the focus is not only on the strongest claim, but on the wider set of stronger claims. Thus, *Match in Order of Strength* might be better aligned with the initial Local Relevance idea, that weak claims can remain part of the overall decision of whom to save, but only when the stronger claims are balanced.

Now let us consider the reasons in favour of *Match in Reverse Order of Strength*. These reasons are almost the exact opposite of those for *Match in Order of Strength*. Firstly, *Match in Reverse Order of Strength*, by matching the weaker claims first in ambiguous cases, leaves the stronger claims in the competition for longer. In contrast to the above, we might think that *Match in Reverse Order of Strength* better respects the wider set of stronger claims (rather than just the strongest claim) as leaving them in the competition for longer increases the likelihood that these claims will be the decisive ones. Again, because of *Match by Closeness*, in practice this does not make a huge amount of difference - in some cases *Match in Reverse Order of Strength* will end up favouring the group

Ambiguity

with the strongest claims and sometimes not, as can be seen in Cases 12 and 13 - but one might still think this logic is preferable.

Secondly, *Match in Reverse Order of Strength* ensures that weaker claims get some consideration, even if taking them into consideration does not affect the outcome. By matching from the bottom up in ambiguous cases, not only do we leave the strongest claims unmatched for longer, but we also give some consideration to the weakest claims in a way that might show a level of procedural respect. Starting from the bottom helps to telegraph that weak claims are not completely irrelevant, even if they cannot in the end play a role in whom we ought to save.

To be completely honest, I find it quite difficult to decide between these two supplementary principles. I cannot see a decisive reason to prefer one over the other, and I can see the logic behind either principle. Perhaps there are decisive reasons in favour of one supplementary principle or another, but if there are I have not found them. Or perhaps this indicates that this is an area of moral free choice, that either way of matching the claims is permissible, even if they lead to saving different groups.

One or the Other

This leads us nicely on to the last way in which we might settle the ambiguity: *One or the Other*. *One or the Other* states that when there is an ambiguity as to how to match claims (and how we settle that ambiguity will affect whom we save²⁹) then neither group is decisively the winner. Of course, we still have to save one of the groups and so perhaps we have a free choice, or perhaps we ought to flip a coin between the two groups. *One or the Other* has the advantage of giving us moral free space over what to do which is an intuitive feature in common sense morality. Either option is permissible, and it is left up to the moral judgement of the decision makers.

²⁹ This clause is not explicit in Van Gils and Tomlin's formulation of the principle. However, there are cases where it might be ambiguous as to how to match claims, but whichever way we choose to settle the ambiguity one group will always win. In such cases, ambiguity should clearly not prevent one group from being decisive.

However, I think flipping a coin here is the wrong approach. We flip a coin when everything is tied because each side is equally weighted. In these cases, it is not that the groups are equally weighted but that it is ambiguous how to match the claims and thus ambiguous whether certain claims are relevant to our decision. It does not seem like we should be flipping a coin to decide whether weaker claims are relevant or not. Doing so seems especially disrespectful. Claims should be deemed irrelevant by the features of the situation - not by chance.

Whilst I am tempted to say that either option is permissible (at the very least I have not resolved whether one option is preferable to the other), I do think there is still good reason to choose one of the two principles and stick to it. If we use *Match in Order of Strength* in some cases and *Match in Reverse Order of Strength* in others, then claimants might object to the principle we choose in their particular case if it does not favour them. Whereas if we choose one principle and stick to it, then this is more easily justified to all the claimants. A further benefit of this approach is that both *Match in Order of Strength* and *Match in Reverse Order of Strength* are more decisive, easier to implement, and are much more amenable to *Sequential Claims-Matching by Closeness*.

So, how will I approach this ambiguity in the rest of the thesis? Well, I mildly prefer *Match in Order of Strength* so shall use this principle throughout the thesis. Firstly, it is important to note that nothing in the rest of the thesis hangs on this decision. If one prefers *Match in Reverse Order of Strength*, then the account I give can be easily amended to use such a principle. In this chapter I developed *Sequential Claims Matching by Closeness* using *Match in Order of Strength*, but we could equally develop it using *Match in Reverse Order of Strength*. To do so, one simply needs to modify Steps III and VI such that we start with the weakest claim. For instance, Step III would become:

- III. Identify the weakest remaining claim *with relevant competing claims one level higher*. Match these claims [with the claims one level higher] and eliminate them from consideration.

Now, my reasons for preferring *Match in Order of Strength* as the supplementary principle are weak but, given the above, they do not need to be strong reasons. I ever so slightly prefer the justification

Alternatives

of, or the logic behind, *Match in Order of Strength*. The justification seems closer to the initial idea behind Local Relevance, in prioritising its focus on the stronger claims. Plus, I think the way in which *Match in Order of Strength* treats claims with equidistant weaker and stronger competing claims is better. This is because it matches such claims with the stronger competing claims, giving the benefit of the doubt to such claims.

Furthermore, there is good practical reason to prefer *Match in Order of Strength*. As we have seen, *Match in Order of Strength* is already endorsed by a number of philosophers in the literature, even if for bad reasons. Thus, *Match in Order of Strength* should be more familiar, and so by using it I should aid understanding and uptake of the rest of the view I develop.

Lastly, we should note that such ambiguities are common in these toy examples, but in reality, such cases will be few and far between. It is unlikely that there will be many claims that we will judge to be exactly equidistant between competing claims of different strength. When there are claims that seem equidistant between competing claims, the first approach should be to identify which of the two comparisons is truly the closer and which is the more respectful matching. Decision makers should use their own moral judgements and only when they really cannot settle between the two options should they use a principle like *Match in Order of Strength* to settle how to compare claims.

Alternatives

Before we conclude this chapter there is one last thing that I ought to address. We might think that the justification behind *Match by Closeness* could also justify a number of other matching principles, and thus subscribing to *Match by Closeness* occurs too quickly here. As such, let us quickly look at some alternatives.³⁰

Now, the idea behind *Match by Closeness* is that it is more disrespectful to compare dissimilar claims than it is to compare similar. Thus, there is a sort of continuity of disrespect. The way this is then

³⁰ Thanks to an anonymous reviewer for encouraging me to address the issue.

fleshed out in *Match by Closeness* is by matching in the most respectful way first and so on.

However, there may be other ways to match with this continuity of disrespect in mind, and we might think that *Match by Closeness* might actually not be the most respectful way of matching claims.

Firstly, look at the following case:

Case 14

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia	2	2
3 - Lost Limb		10
4 - Lost Hand		

In this case, *Match by Closeness* tells us to match the paraplegia claims together first and then to compare the death and lost limb claims. Now, whilst matching this way does provide the most respectful matching possible to the paraplegia claimants - their claims are matched against other claimants of the same strength - it also ensures a more disrespectful matching for the death and lost limb claimants - their claims are matched over two levels rather than one.

A different matching principle could use the same justification of reducing disrespectful comparisons, to say that we ought to compare the A1 claims with the B2 claims and the A2 claims with the B3 claims, such that all claims are compared with claims of equal distance. Thus, we might think that *Match by Closeness* fails to actually fulfil its stated aim of reducing disrespectful comparisons.

Firstly, whilst there is something to this criticism of *Match by Closeness*, I want to show why I do not think this undermines the principle. If we were solely focused on minimising the amount of total disrespect in our matching of claims, then I think this objection would have more force. However, this is not the only consideration we have. As we noted above, we also want to respect the strongest claims and give them priority. As such, we have a bit of a trade off in these cases. If the strongest claimant is not to be saved, then this claimant will want her claim to be matched in the most

Alternatives

respectful way - i.e., to the closest competing claim. However, the strongest claimant will also want their claim to remain in the competition for as long as possible such that it can maximise its chances of being the decisive claim.³¹

I think *Match by Closeness* captures a nice balance here. By matching the A2 and B2 claims together, it does mean in a case like Case 14 that the A1 claims are compared with the more distant B3 claims. But it also means that in other cases - like the First Horn cases above, when instead of B3 claims there are B4 claims - the A1 claims can make the B4 claims irrelevant and thus be decisive. In this way the strongest claimants should prefer the risk of a less respectful comparison, as it gives them a greater chance of a positive outcome. They trade procedural respect for the chance of a better outcome.

With my defence of *Match by Closeness* out of the way, let us also briefly look at why the other alternatives would fail. Here are several alternatives:

Lexically Match by Closeness: first match the strongest claims with closest relevant claims, then the next strongest claims, and so on.

Minimise Total Distance: when matching claims, match claims such that the total distance between compared claims is minimised.

Minimise Distant Matchings: when matching claims, match claims such that the number of distant matchings is minimised.

Firstly, I include *Lexically Match by Closeness* for completeness, but we should dismiss this option from the get-go. This approach is actually just a less ambiguous version of *Match in Order of Strength* from the last chapter. We saw then that such a view lands on the First Horn of Horton's dilemma and violates several important principles.

³¹ In fact, weaker claims will also have this tension. They too will want to have their claims matched in the most respectful way, but also to have their claims in the competition for as long as possible such that they might be decisive.

Secondly, let us look at *Minimise Total Distance*. This principle sticks closely to the reasoning behind *Match by Closeness*, perhaps better than *Match by Closeness* itself. Rather than say we ought to match claims in order of respect, *Minimise Total Distance* tells us that we ought to consider the total disrespect generated by our matchings. Thus, it might be better to match some claims in less respectful ways, if it means we can match other claims in more respectful ways. For instance, in Case 14, we match the paraplegia claims in a less respectful manner, such that we can match the death claim more respectfully. In this way we might think of *Minimise Total Distance* as consequentialising the respect element of the limited aggregation view.

However, *Minimise Total Distance* has a couple of issues. Firstly, it is not entirely clear how we would calculate the total distance between matched claims, or the total amount of disrespect generated by such matchings. How do we aggregate distance or disrespect? Are two one-level comparisons equivalent to one two-level and one zero-level comparison? Secondly, even if we can settle this ambiguity, *Minimise Total Distance* also lands on the fatal First Horn of Horton’s dilemma:

First Horn Sharpened

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia	(100)	
3 - Lost Limb		(6)
4 - Lost Hand		1000

Of course, at Stage 1 we ought to save Group A. However, *Minimise Total Distance* will tell us at Stage 2, that we minimise the total distance of compared claims by comparing the A1 claim with the B3 claims, and the A2 claims with the B4 claims. Thus, despite Group A being strengthened by more and stronger claims, *Minimise Total Distance* tells us to swap to saving Group B. As we have already seen, this violates *Greater Consideration for Stronger Claims*, and the *Principle of Net Addition*.

Lastly, *Minimise Distant Matchings*. This principle is less ambiguous than *Minimise Total Distance* as there is no aggregation or calculation for total distance that we need to work out. *Minimise Distant Matchings* is closer to *Match by Closeness* in the sense that it focuses on particular matchings rather

Concluding Remarks

than total matchings. However, rather than maximise the number of close matchings like *Match by Closeness* does, *Minimise Distant Matchings* minimises the number of distant matchings. They are concerned with different ends of the spectrum - analogous to the difference between utilitarianism and negative-utilitarianism.

However, *Minimise Distant Matchings* also lands on the First Horn of Horton's dilemma. In *First Horn Sharpened*, *Minimise Distant Matchings* tells us to save Group B at Stage 2. This is because we ought to minimise distant comparisons, and so we avoid comparing the A1 claim with the B4 claims. Thus, we match the A1 claim with the B3 claims, the A2 claims with the B4 claims and save Group B.

Therefore, none of these alternative views can capture what we want to say about limited aggregation, even if they have some claim to the same justification behind *Match by Closeness*.

Concluding Remarks

To conclude, *Match by Closeness* (supplemented by *Match in Order of Strength*) gives us a decisive way to match claims, and I have formulated this into *Sequential Claims Matching by Closeness*. It has a strong independent rationale based on matching claims in the most respectful way possible: we compare equal claims first, then similar claims, and then more distant claims, such that only if the earlier steps are not decisive do we weigh strong claims against weak. This tracks with the underlying justification of most limited aggregation views - namely that we should avoid comparing weak and strong claims as doing so does not respect the partial concern of each claimant and what is at stake for them - and is thus amenable to these relevance accounts.

We have also seen how *Match by Closeness* respects the special importance of the strongest claim in a way that is not detrimental to that claim. We have seen how it avoids the First Horn of Horton's dilemma, and thus how it respects *Equal Consideration for Equal Claims*, *Greater Consideration for Stronger Claims*, and the *Principle of Net Addition*. Plus, we saw how this principle can give us

independently good reason to swap whom we save in the Second Horn cases, by changing the moral shape of such cases, and bringing to the forefront more respectful comparisons.

We also looked at some remaining ambiguities for my view. Namely, that *Match by Closeness* cannot tell us the order in which we ought to save claims when those claims are equidistant. I showed that supplementing the view with either *Match in Order of Strength* or *Match in Reverse Order of Strength* could settle this ambiguity. I argued that, even though they give different answers in different cases, we have a free choice as to which of these two principles to use, and that this gives the view a valuable moral free space. Nevertheless, I decided to settle on *Match in Order of Strength* as the supplementary principle which we will use to settle ambiguities going forward.

Lastly, we considered some alternative accounts that might also capture the reasoning behind *Match by Closeness*. We saw that such accounts all fatally land on the First Horn of Horton's dilemma and thus violate *Equal Consideration for Equal Claims*, *Greater Consideration for Stronger Claims*, and the *Principle of Net Addition*. However, by looking at these alternative accounts we were able to better recognise a tension within limited aggregation - namely that there is a trade-off between comparing claims in the most respectful way possible and keeping the strongest claims in the competition such that they have a greater chance of being the decisive claims. This further strengthened the *Match by Closeness* view, as we can see how it finds a good balance between these differing considerations.

However, *Match by Closeness*, like Van Gils and Tomlin's view (but unlike Steuwer's), cannot currently capture Kamm's sore throat cases (Kamm, 1998, p. 101). I will turn my focus to this issue in the second half of the thesis.

Part 2 – Tiebreaking Claims

Chapter 4 – Tiebreaks and Two Types of Relevance

Introduction³²

Relevance views come in two flavours: Local Relevance and Global Relevance. In Chapter 1, I argue that Horton's dilemma means we ought to reject all Global Relevance views, but in this chapter I will present a common trilemma for both views. Namely, that neither view can capture our intuition in tiebreak cases, without forfeiting our intuitions in other important cases.

The chapter then presents a way to salvage relevance views and capture all our intuitions using a Hybrid view. By distinguishing between two types of relevance we can combine the strengths of Local and Global Relevance views such that we can hold all our intuitions, consistently and in a non-ad-hoc manner. Building on this, the chapter demonstrates how we might amend the strongest formulation of a Relevance view, into a Hybrid account.

The main focus of this chapter is to develop a relevance view that can capture all the intuitions we have; however, deeper justifications will still be needed for a full account of any relevance view. Thus, the end of the chapter considers what deeper justification might support the Hybrid view, indicating the direction such literature might go.

³² A version of this chapter has been published in *Ethical Theory and Moral Practice* and is available online at: <http://dx.doi.org/10.1007/s10677-022-10270-3>.

We noted at the end of the last chapter that the view so far developed in this thesis cannot capture Kamm's sore throat intuition:

Kamm's Sore Throat Case: We can save either person A from death, or person B from death and cure person C's sore throat.

It seems that a small claim, such as a sore throat claim, should not be the reason that A does not have a chance to be saved. What we ought to do in Kamm's Sore Throat Case is to flip a coin between saving A and saving B. That way the two dying people have an equal opportunity to be saved. Local Relevance views, however, have a problem with capturing this intuition:

According to Local Relevance, C's sore throat cannot outweigh A's death, but it may be relevant more generally to what I am permitted or required to do. A's claim to be cured from the lethal illness is exactly counterbalanced by B's identical claim. Therefore, there is nothing left for C's claim to outweigh. Therefore, Local Relevance implies that I ought to cure B and C rather than flipping a coin, which seems wrong. (Tadros, 2019, pp. 184-185)

Some limited aggregationists are willing to bite the bullet on this as the "the smallest caliber bullet" (Van Gils & Tomlin, 2019, p. 253) (Rüger, 2020, pp. 464-465) and some may not share Kamm's intuition at all. However, I find Kamm's case compelling. Further examples are more powerful still. It seems clear that neither a single speck of dust in an eye nor a fleeting itch should decide between two lives. Thus, any theory that can capture these intuitions has a significant prima facie advantage over one that cannot.

The rest of this chapter will be spent exploring how we might amend my view, and wider Relevance views in general, to accommodate the Sore Throat Intuition. By the end of the chapter, I will have demonstrated that the only way to accommodate them is to bring back Global Relevance and combine the insights of both types of relevance into a Hybrid view.

Global and Local Relevance

Before we jump into the problem it would be helpful to refamiliarise ourselves with the distinction between Global and Local Relevance, and why we adopted Local Relevance in the first place. The following are two simple formulations of Global and Local Relevance.

Global Relevance: A claim, X, is Globally Relevant if X is strong enough compared to the strongest claim with which it competes.

Local Relevance: A claim, X, is Locally Relevant to another claim, Y, if X is strong enough compared to Y.

Thus, we can see that Global Relevance views decide whether a claim is relevant or irrelevant at the outset, and that relevance does not change. They do so by examining whether a claim is strong enough compared to the strongest competing claim. If a claim is not strong enough, then it can play no part in which group to save, even when the strongest competing claims is matched.

On the other hand, Local Relevance views consider whether a claim is relevant or irrelevant 'on the go', and the claim's relevance might change depending on how other claims are matched. Local Relevance views determine relevance by examining whether a claim is strong enough to *each and every* type of claim with which it competes to be able to match with that claim. So, a claim might be relevant to a weak claim in the competition, but irrelevant to a stronger competing claim. Because of this a claim that is irrelevant to the strongest claim can still play a role in deciding which group to save.

Thus, the two views handle Kamm's Sore Throat case very differently. Global Relevance views can capture the sore throat intuition. Such views tell us that the sore throat claim is irrelevant to the death claims and should thus be dismissed from consideration at the outset. Given that everything else balances, Global Relevance will tell us to flip a coin. On the other hand, Local Relevance, as we

have already seen, will allow the sore throat claim to remain part of the overall decision when all the other features are matched and will thus tell us to save B and C outright.

We might think that this gives us good reason to prefer the Global Relevance approach. However, the Global Relevance approach is fatally flawed. Firstly, Patrick Tomlin identifies a number of powerful counterexamples to the Global Relevance approach (Tomlin, 2017), which we will not go into here. Secondly, as we saw in Chapter 1, the Global Relevance approach lands on the fatal First Horn of Joe Horton’s ‘Fatal Dilemma’ (Horton, 2018). Moreover, Global Relevance struggles to capture our intuitions in other tie breaking cases. Consider the following:

Lost Hand Tiebreakers. We can save Group A or Group B. In Group A we have a death claim, in Group B we have a death claim and a hundred lost hand claims.

Let us suppose that our respect-based reasons tell us that lost hand claims are too weak to be relevant to death claims - no number of lost hand claims should ever be saved instead of a life. In this case, Global Relevance will tell us to flip a coin. This is because the lost hand claims are irrelevant to the death claim with which they compete, and therefore cannot play any part in deciding who wins.

But this seems wrong. It seems that we should not completely ignore the lost hand claims in this case. The lost hand claims, whilst not strong enough to compete and outweigh a death claim, are strong enough to choose between two differing death claims. It seems disrespectful to the people whose hands are in danger to ignore their claims when all else is equal. Lost hands are unlike sore throats in this way. Sore throats seem to lose all moral relevance when compared to what is at stake for a dying person, but lost hands do not.³³ I will return to the question of why this might be in the Deeper Justification section.

³³ To put it in Voorhoeve’s terms, the partial concern that lost hand claimants have for their claim does not permit them to prefer their own rescue over the rescue of a dying person, but partial concern does permit them to prefer their own rescue over an equal opportunity of rescue for a competing death claim.

Introduction

Firstly though, the intuition that *Lost Hand Tiebreakers* gets at might not be immediately clear. So, let us test it. Suppose you are a doctor deciding whether to send your limited resources to Ward A or Ward B. In Ward A there is a dying man, and in Ward B there are a hundred people with strong, yet (compared to death) irrelevant claims. To reiterate, in this thesis I use lost hand claims, but they could be lost finger claims, lost limb claims or some other strong claim that is slightly too weak to ever aggregate against a death claim.³⁴ Let us call this:

Death versus Lost Hands. We can save Ward A or Ward B. In Ward A we have a death claim, in Ward B we have a hundred lost hand claims.

Suppose that after some deliberation, you work out that the lost hand claims are irrelevant and so send the resources to Ward A. However, if the hundred people in Ward B had slightly stronger claims, then their claims would have been relevant and by sheer weight of number significantly outweighed the death claim. Now, just before you send off those resources to Ward A, you are informed that a dying person has been wheeled into Ward B, and you could save her life as well as the hundred. Whom should you send your resources to? My strong intuition is that we ought to send the resources to Ward B rather than flip a coin. Whilst the hundred do not have claims strong enough to aggregate against the dying person, they do have claims strong enough to be decisive in a tiebreak situation. In fact, it seems that it would be disrespectful to not include the lost hand claims in our decision when all other features are balanced. Global Relevance cannot capture this intuition - on any Global Relevance view claims are either always or never taken into consideration, there is no gradation.

Now, Local Relevance solves this problem. According to Local Relevance, in *Death versus Lost Hands*, the lost hand claims are irrelevant to the death claim and therefore we must save Group A. But in *Lost Hand Tiebreakers*, the lost hand claims, whilst irrelevant to the death claims, can remain

³⁴ To reiterate, readers are encouraged to replace lost hand claims throughout the thesis with the strongest claims they can imagine that nevertheless remain irrelevant to claims against death.

relevant to what we ought to do all things considered. Because the death claims counterbalance each other, the lost hand claims are not thrown into direct comparison with the death claims. Thus, we ought to take the lost hand claims into consideration, and given that everything else is equal we should save Group B. This is a much better result.

A Trilemma

Given the above, we might think that we just need to amend Local Relevance, perhaps add an additional clause, to handle *Kamm’s Sore Throat Case*. However, as we will see, no amendment can capture our intuition in this case without sacrificing our intuitions in other cases. In fact, no single relevance view can capture all our intuitions. To see why consider a trilemma generated by *Kamm’s Sore Throat Case*, *Lost Hand Tiebreakers* and *Death versus Lost Hands*:

Kamm’s Sore Throat Case

Group A	Group B
1 person facing death	1 person facing death 1 person facing a sore throat

Lost Hand Tiebreakers

Group A	Group B
1 person facing death	1 person facing death 100 people facing a lost hand

Death versus Lost Hands

Group A	Group B
1 person facing death	100 people facing a lost hand

Intuitively, we should flip a coin in *Kamm’s Sore Throat Case*, save Group B in *Lost Hand Tiebreakers* and save Group A in *Death versus Lost Hands*. Yet no single relevance view can capture what we want to say in all three of these cases.

Global Relevance gets the right answer in *Kamm’s Sore Throat Case* and in *Death versus Lost Hands*. But it tells us to flip a coin in *Lost Hand Tiebreakers*, when it seems we ought to save Group B. Local

A Trilemma

Relevance gets the right answer in *Lost Hand Tiebreakers* and *Death versus Lost Hands*. But it tells us to save Group B rather than flip a coin in *Kamm's Sore Throat Case*.

Now let us consider amendments to the Local Relevance view. Victor Tadros argues for the following:

if I choose to save B because I can cure C's sore throat, C's sore throat grounds a permission not to save A from death where otherwise I would be required to give him an equal chance of being saved. Friends of Local Relevance might claim that this is objectionable—C's sore throat cannot ground a permission not to give A a chance of being saved from death, either on its own, or in conjunction with other permission-grounding types of facts. It can only outweigh duty-grounding types of facts that do not disable its duty-grounding force. (Tadros, 2019, p. 185)

But if Tadros takes this position, he is unable to capture our intuition in one of either *Lost Hand Tiebreakers* or *Death versus Lost Hands*. Suppose lost hands can ground a permission not to give another claimant a chance of being saved from death - such that the lost hand claims are relevant. If so, Tadros' view will get the wrong answer in *Death versus Lost Hands* - the lost hand claims would be able to outweigh the death claim. Conversely, if lost hands cannot ground such a permission, then in *Lost Hand Tiebreakers* the lost hand claims cannot be decisive in favour of Group B, again getting the wrong answer.³⁵

Similarly, Van Gils and Tomlin suggest a possible sore throat amendment to their version of Local Relevance, the Sequential Claims Matching view. To understand this amendment, we need to understand the role of anchoring claims. Anchoring claims are not just a feature of their account and they will arise throughout the thesis, so it is important we understand them properly.

³⁵ If Tadros were to recognise that the strength a claim needs to be relevant to a chance to be saved differed from the strength needed to be relevant to a claim to be saved, then he could capture all these cases. To do so he would need to distinguish between two types of relevance, in the same way that I do in the next section. Still, Tadros' approach here helpfully supports the deeper justification for my view.

The anchoring claim is the strongest claim under consideration (or claims if there are multiple claims of the same strength); in the cases above, the death claims are the original anchoring claims. It is the anchoring claim we compare other claims to, to determine whether they are sufficiently strong to be relevant. In this way the person with the anchoring claim has a pro tanto claim to be saved - unless the claim is silenced or outweighed by other claims relevant to it, we must save the anchoring claimant. On a Local Relevance view, when an anchoring claim is matched, the strongest unmatched claim becomes the new anchoring claim. This is how claims that were not relevant to the original anchoring claim might still be relevant to the overall decision by being relevant to new anchoring claims.

Van Gils and Tomlin then suggest the following amendment to their account:

We could alter Sequential Claims-Matching so that when it comes to considering ‘new anchoring claims’, we only consider those claims which are relevant to the weakest of those claims that have been previously matched and set aside. (Van Gils & Tomlin, 2019, p. 233)

This amendment allows their version of Local Relevance to capture our intuition in *Kamm’s Sore Throat Case*. Because the sore throat is not relevant to the weakest previously matched claim (i.e., the death claim) it cannot become the new anchoring claim and so we must flip a coin. Naturally, however, this amendment also loses our intuition in *Lost Hand Tiebreakers*. Before we could say that, whilst the lost hand claims were irrelevant to the death claims, they could still play a role in our deliberation. However, with the sore throat amendment, the lost hand claims can no longer play a role as they are not relevant to the weakest previously matched claim.

So, no account so far considered can capture all these intuitions about relevance. Perhaps, there are good reasons for biting the bullet in one or more of these cases. However, a view that could capture our intuition in all three cases would be prima facie stronger and deserves serious consideration.

Two Types of Relevance

The problem above seems to be that there is a tension as to how far we want relevance to extend. In *Death versus Lost Hands* we want to say that only very similar claims are relevant, whilst in *Lost Hand Tiebreakers* we want to say that more distant claims are relevant (whilst still allowing even more distant claims to remain irrelevant). We are being pulled in two directions.

This tension suggests that what we have previously treated as a single phenomenon might in fact be two. We have been wrong to consider only one type of relevance; instead we need to differentiate between types of relevance more finely. We will explore the deeper rationale behind such a distinction in a later section. First let us consider how exactly we might distinguish between types of relevance, and whether the distinction allows us to untangle the tension between these cases.

The tension between *Lost Hand Tiebreakers* and *Death versus Lost Hands* is that relevance is needed to play two distinct roles in the different cases. In *Death versus Lost Hands* relevance is used to determine whether a claim can compete with another claim, i.e., whether the hundred lost hand claims can match with the death claim. In *Lost Hand Tiebreakers*, relevance is used to determine whether a claim can become, what Van Gils and Tomlin called, the anchoring claim, i.e., whether the lost hand claims can be decisive between two death claims. These two roles we might call the Matching Role and the Anchoring Role:

Matching Role: the role relevance plays in determining whether a claim can match or outweigh a competing claim.

Anchoring Role: the role relevance plays in determining whether a claim can become the new anchoring claim.

Up until now, we have had only one type of relevance to play both these roles, and this is what leads us to the tension in the view. If we want to capture all our intuitions, then we need to separate these

roles from each other and distinguish between two types of relevance that can play the different roles.

Fortuitously, we have already distinguished between two types of relevance - Global Relevance and Local Relevance. Let us reintroduce the two types of relevance here:

Global Relevance: A claim, X, is Globally Relevant if X is strong enough compared to the strongest claim with which it competes.

Local Relevance: A claim, X, is Locally Relevant to another claim, Y, if X is strong enough compared to Y.

With this in view, we can see that there are four notions of relevance in play. We have Global Relevance regarding the Matching Role, Global Relevance regarding the Anchoring Role, Local Relevance regarding the Matching Role, and Local Relevance regarding the Anchoring Role.

Global Relevance and Local Relevance have previously been considered rival understandings of relevance, with proponents of either view ascribing both the Matching and Anchoring roles to their preferred type of relevance. However, there is no reason to suppose we must choose one form of relevance for both roles. Instead, we could choose one type of relevance for one role, and the other type of relevance for the other role - I call this the Hybrid Approach.³⁶

In the light of this, Local Relevance seems most appropriate for the Matching Role, and Global Relevance most appropriate for the Anchoring Role. Let us quickly consider why.

If we used Global Relevance for the Matching Role, whether a claim can be aggregated against a competing claim would not be determined by whether it is relevant to that particular claim, but

³⁶ I should note here that I am not the first to attempt a hybrid relevance view. Bastian Steuwer has very recently offered his own hybrid view, *Hybrid Balance Relevant Claims*, that also attempts to avoid the pitfalls of both the Global and Local Relevance views. However, as we have seen in Chapter 2, this hybrid view has a number of fatal flaws. It will also not be able to escape the Trilemma identified above - *Hybrid Balance Relevant Claims* will tell us to flip a coin in *Lost Hand Tiebreakers* - as Steuwer does not distinguish between the Matching and Anchoring Roles.

Two Types of Relevance

whether it is relevant to the strongest claim in the competition. This seems overzealous - the strongest claim in the competition should not determine whether two smaller sets of claims are comparable! However, if we use Local Relevance, we determine whether a claim can be aggregated against a competing claim, by looking to see if it is relevant to that particular claim. This seems much more intuitive.

Where the Anchoring Role determines whether a claim can become the new anchoring claim, it also determines whether the last remaining claim can tiebreak between the already matched claims. It seems obvious that Global Relevance is more appropriate for this role. If a claim is going to tiebreak between claims, not only is it tiebreaking between weaker claims, but it is also, of course, tiebreaking between the strongest claims. Thus, whether it is strong enough to become the new anchoring claim should be determined in comparison with the strongest overall claim.

I posit we formalise this by combining Local Relevance and the Matching Role into Competing Relevance, and Global Relevance and the Anchoring Role into Anchoring Relevance:

Competing Relevance: the relevance a claim needs to be able to match or outweigh a competing claim. A claim, X, has competing relevance with another claim, Y, if X is sufficiently close in strength to Y such that X can match with Y.

Anchoring Relevance: the relevance a claim needs to be able to become the new anchoring claim. A claim, X, has anchoring relevance if X is sufficiently close in strength to the strongest claim (matched or unmatched) with which X is in competition, such that X can become the new anchoring claim when all stronger claims are matched.

Once we recognise this distinction, we can set Competing Relevance and Anchoring Relevance to extend different distances. For instance, we can set Competing Relevance to extend only to very similar claims, such that no claim weaker than a lost limb claim can match with a death claim. We

can do this, whilst simultaneously setting Anchoring Relevance to extend further, such that lost hand claims can become the new anchoring claims when all death claims are matched.

Consider how this might work in *Kamm’s Sore Throat Case*, *Lost Hand Tiebreakers* and *Death versus*

Lost Hands:

Kamm’s Sore Throat Case

Group A	Group B
1 person facing death	1 person facing death 1 person facing a sore throat

Lost Hand Tiebreakers

Group A	Group B
1 person facing death	1 person facing death 100 people facing a lost hand

Death versus Lost Hands

Group A	Group B
1 person facing death	100 people facing a lost hand

With Competing Relevance extending only as far as lost limb claims, in *Death versus Lost Hands* the hundred lost hand claims cannot be matched against the death claim. Thus, Group A wins in *Death versus Lost Hands*. Simultaneously, with Anchoring Relevance extending further, in *Lost Hand Tiebreakers* the lost hand claims can become the new anchoring claims, and therefore be decisive in favour of Group B. Yet, Anchoring Relevance would not extend so far as to include sore throats and so we must flip a coin in *Kamm’s Sore Throat Case*. So, we get the right answer in all three cases.

Amending Sequential Claims Matching by Closeness

So far, I have spoken in generalities about Relevance views and how the Hybrid Approach might be used. But now I want to turn to how we can combine this Hybrid Approach with the *Sequential Claims Matching by Closeness* view developed in the last chapter. This will significantly strengthen

Amending Sequential Claims Matching by Closeness

the view and it will work as a proof of concept for how others might adapt their Relevance views to include the Hybrid Approach.

In *Sequential Claims Matching by Closeness*, we have a procedure which compares competing claims according to their closeness. First, we compare identical claims, then relevantly similar claims, and then wider relevant claims, until we reach the limit of relevance. *Sequential Claims Matching by Closeness* does so by splitting claims into 'levels' of claims sorted by their strength, with the relative strength of claims between levels remaining the same. Let Level 1 Claims be death claims, Level 2 Claims be paraplegia claims, Level 3 Claims be lost limb claims and so on. Remember, that each level has half the weight of the next level up, such that two Level 2 claims, match with exactly one Level 1 claim, and four Level 3 claims match with exactly two Level 2 claims or one Level 1 claim.

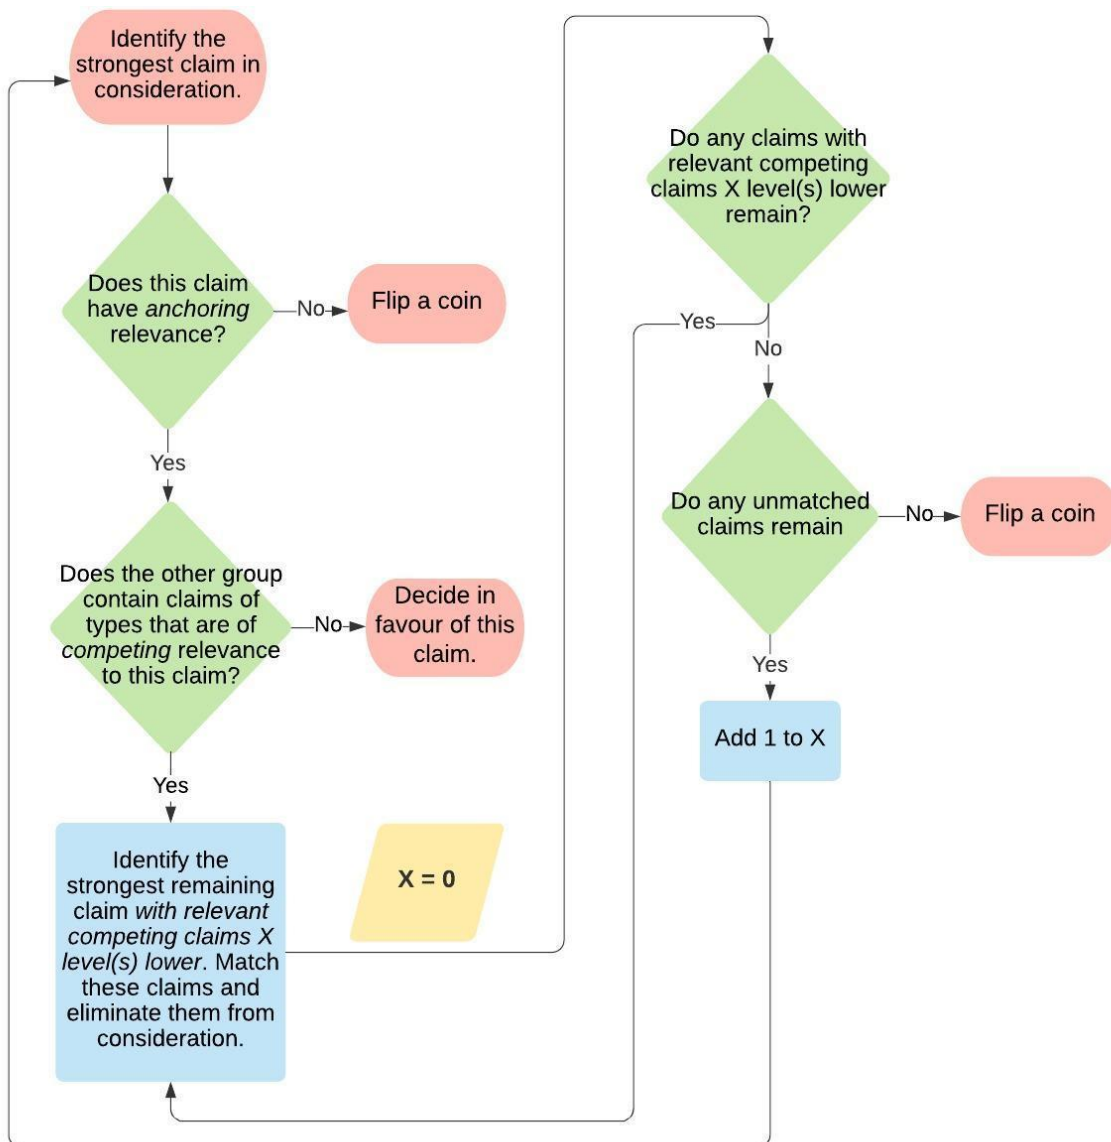
The following is an amended version of *Sequential Claims Matching by Closeness* which distinguishes between the two types of relevance, to capture the Hybrid Approach. Again, I include a flow chart to help illuminate the sequence.

Sequential Claims Matching by Closeness

- I. Identify the strongest claim still in consideration and ask, does it have *anchoring* relevance?
 - a. If so, go to Step II.
 - b. If not, then it is not the case that you should decide in favour of one group over the other (though you must save one).
- II. Does the other group contain claims of types that are of *competing* relevance to this claim and sufficient in number to match it?
 - a. If so, go to Step III.
 - b. If not, decide in favour of the group with the strongest claim still in consideration.
- III. Identify the strongest remaining claim *with* competing claims [zero] levels lower. Match these claims and eliminate them from consideration.
- IV. Continue step III until there are no claims with competing claims [zero] levels lower.

- V. Repeat steps I - IV, with competing claims one level lower, then two levels lower, and so on.
- VI. Continue until either:
 - a. One group contains unmatched claims with anchoring relevance, 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).

Sequential Claims-Matching by Closeness – Flow Chart



Amending Sequential Claims Matching by Closeness

Once we adapt the sequence in this way, we fully capture the two types of relevance within the sequence. Therefore, we can redeem the sore throat intuition, whilst also getting the right answer in *Lost Hand Tiebreakers* and *Death versus Lost Hands*. Consider again *Kamm's Sore Throat Case*:

Kamm's Sore Throat Case

Level	Group A	Group B
1 - Death	1	1
2 - Paraplegia		
3 - Lost Limb		
4 - Lost Hand		
...		
100 - Sore throat		1

Remember from previous chapters that the levels refer to claims against death, paraplegia, lost limbs, lost hands, and sore throats, respectively and that, for ease of reference, we will call the Level 1 claims in Groups A and B, the A1 and B1 claims, and the Level 100 claim in Group B, the B100 claim.

Now we can set the extent of relevance. For the purposes of this chapter, let us say that Competing Relevance is set at two levels, such that Level 3 claims have Competing Relevance to Level 1 claims, but Level 4 and below claims do not. Similarly, let us set the extent of Anchoring Relevance at four levels, such that Level 5 claims have Anchoring Relevance to Level 1 claims, but Level 6 and below claims do not. With this in view, let us go through *Kamm's Sore Throat Case* slowly to demonstrate how the sequence works:

Step I: First we identify the strongest claim, which is A1 (or B1). As the overall strongest claim in the competition it, trivially, has Anchoring Relevance.

Step II: We check if the A1 claim has claims of Competing Relevance to it and whether these claims can match it - it does with the B1 claim, so we move to Step III.

Chapter 4 – Tiebreaks and Two Types of Relevance

Step III: We identify the strongest claim with claims zero levels lower (i.e., claims on the same level). Again, this is the A1 claim, and we match it with the B1 claim. I have demonstrated this on the below table by crossing through the A1 and B1 Claim.

Kamm's Sore Throat Case

Level	Group A	Group B
1 - Death	1	1
2 - Paraplegia		
3 - Lost Limb		
4 - Lost Hand		
...		
100 - Sore throat		1

Step IV and V: We check to see if there are any unmatched claims, which there are so we loop back round to Step I.

Step I (Round 2): On the second pass through the sequence, we identify the B100 claim as the strongest claim still in consideration, we check to see if it has Anchoring Relevance.

Anchoring Relevance only extends four levels, and thus we can see that the B100 claim does not have Anchoring Relevance, and so we must flip a coin.

Thus, by amending the sequence in this way we have managed to capture the sore throat intuition!

Let us also, quickly consider *Lost Hand Tiebreakers* and *Death versus Lost Hands*:

Lost Hand Tiebreakers

Level	Group A	Group B
1 - Death	1	1
2 - Paraplegia		
3 - Lost Limb		
4 - Lost Hand		100

On the first pass through the sequence, we match the A1 claim and the B1 claim. On the second pass, at Step I, we identify the B4 claims as the strongest remaining claims, and we check to see if they have anchoring relevance - they do, so we move onto Step II. At Step II, we look to see if there

Amending Sequential Claims Matching by Closeness

are any remaining claims with competing relevance to the B4 claims and, since there are none, Step II tells us to choose in favour of Group B. Our intuition is captured here too.

Death versus Lost Hands

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia		
3 - Lost Limb		
4 - Lost Hand		100

Now *Death versus Lost Hands*. At Step I we identify the A1 claim as the strongest claim, recognise that it has anchoring relevance and move on to Step II. At Step II, we look to see if there are any claims with competing relevance to the A1 claim. The B4 claims are too distant to have Competing Relevance, and so Step II tells us to decide in favour of Group A. Our intuition is captured once again!

A further positive of fleshing out the Anchoring and Competing Relevance distinction in this way is that doing so allows further gradation. A claim without anchoring relevance might not be able to be decisive in favour of a Group, but it can still play a role in our overall decision by preventing other claims from becoming decisive. Consider the following case:

Balanced Tiebreakers

Level	Group A	Group B
1 - Death	1	1
2 - Paraplegia		
3 - Lost Limb		
4 - Lost Hand		
5 - Lost Finger	10	
6 - Broken Arm		30

In this case, anchoring relevance is set at four levels. Thus, the ten A5 claims in Group A do have anchoring relevance, but the thirty B6 claims do not. Nonetheless, the B6 claims can still play a role in the outcome. This is because they still have competing relevance with the A5 claims, and thus can counterbalance and prevent them from being decisive. This means we do not leap straight from claims with relevance to those with no relevance but have a more natural weakening of relevance.

Deeper Justification

Now the primary purpose of this chapter is to vindicate limited aggregation intuitions and defend limited aggregation against the charge of inconsistency. As I have shown, the Hybrid Approach is fully able to capture our intuitions and thus the main task of the chapter is complete.

What is novel about my account is the way in which I apply the different types of relevance to the different roles. I do not introduce any distinction in relevance that has not been previously discussed, I make use of two already plausible conceptions of relevance and simply show that they are compatible. The distinction I do introduce, between the Matching and Anchoring role, simply describes how we already use relevance in two subtly different ways.

As such, there seems to be little reason to think that both roles must be fulfilled by one type of relevance. Hence, no special or extra explanation is needed as to why we should adopt the Hybrid Approach as compared to a non-hybrid approach. That it gets the right answers in the above cases is itself a significant part of the reason to prefer this account over either Local or Global Relevance by itself.

Nonetheless, any limited aggregation account is going to need a deeper theoretical justification if it is to provide a complete explanation, and this account is no exception. Thus, it would be prudent if we spent a little time indicating what theoretical grounding we might have for distinguishing between anchoring and competing relevance.

Firstly, we should note that the way reasons work on this view is not completely unique, and that we already have a framework for explaining why there is an asymmetry in how some of the claims are treated. Jeff McMahan introduces a distinction between reason-giving and cancelling weight when discussing the morality of causing people to exist (McMahan, 2009). He argues that certain goods are reason-giving whilst other goods only cancel, or outweigh, the reasons of the reason-giving goods. Negative value in a life provides reason not to bring that life into existence. Positive value in a

Deeper Justification

life, however, does not provide reason for bringing that life into existence. Nevertheless, that positive value does weigh against the negative value's reason not to bring that life into existence.

Anchoring and competing relevance can be understood as making use of this distinction. Claims with anchoring relevance have reason-giving weight; such claims can provide reasons to save their group. Claims without anchoring relevance do not have this reason-giving weight. However, they might still have cancelling weight. This cancelling weight permits them to outweigh claims with reason-giving weight, but not to provide reason-giving weight of their own. With these two types of weight in view we can see why we also would need two types of relevance. We will need one type of relevance to determine whether a claim has reason-giving weight, and another to determine whether a claim has cancelling weight. It also seems that whether a claim has reason-giving weight ought to be determined by some sort of global relevance condition, whilst whether a claim has cancelling weight need only be determined by being compared to exactly that which it is cancelling.

However, this only gives us high-level theoretical reasons for the distinction. To provide a complete explanation of the account we also need deeper moral explanations for the distinction. Again, to demonstrate the generality of my account I will couch this discussion in the most basic terms of respect and allow readers to insert their preferred rationale behind relevance to fill in the detail.

Here I think two distinctions might be helpful in explaining the two types of relevance:

1. Act disrespect vs procedural disrespect
2. Claims to be saved vs claims to an equal (or proportional) opportunity to be saved

Let us handle the simpler distinction between act disrespect and procedural disrespect, and how this distinction might elucidate the two types of relevance, first. Act respect is about whether an action is respectful or not. For instance, if I choose to save hands over a life, that act is disrespectful because the decision to choose to save hands rather than a life is disrespectful: it is disrespectful to attend to hands when someone's life is on the line, as it fails to recognise the strength of the dying person's

claim, with the strength corresponding to what is at stake for them. On the other hand, if I am faced with a choice between saving two lives, then, whichever life I save, the action is respectful.

Procedural respect, on the other hand, is about the process that leads to the action. For instance, choosing between two life claims by flipping a coin shows procedural respect, because doing so recognises what is at stake for both claimants and gives them an equal chance of being saved. On the other hand, if I choose to save the person with blue eyes then I am showing procedural disrespect even though the act of saving that claimant's life is not itself disrespectful. Similarly, if we were choosing between a life claim and a lost hand claim, and I choose to save the life because the dying person had blue eyes, I would be equally procedurally disrespectful whilst doing the respectful act.

We can now see how this distinction might explain the two types of relevance. Competing Relevance is concerned with action respect. Competing Relevance rules out comparisons between competing claims of different strengths, if such a comparison could lead to a disrespectful action. For instance, it would be a disrespectful action to save a multitude of claimants from lost hands instead of saving a single life, and so we must not match lost hand claims with death claims.

Anchoring Relevance on the other hand is concerned with procedural respect. Certain claims are so weak that it would be disrespectful to even take them into consideration when deciding between stronger claims. Thus, we dismiss considerations such as eye colour, but we will also dismiss very small considerations such as sore throats. In the same way that it is disrespectful to choose between two death claims because of eye colour it is disrespectful to choose between two death claims because of a sore throat.

The second distinction makes use of the difference between a claimant's claim to be saved, and a claimant's claim to an equal (or proportional) opportunity to be saved.³⁷ It is important to note that

³⁷ As we have seen above, Tadros says something like this (2019, pp. 185-186).

Deeper Justification

a claimant's claim to be saved from X is significantly stronger than another claimant's claim to an equal opportunity to be saved from X. This is why a claim to be saved from death will outweigh a claim to an equal opportunity to be saved from death. It is why we should choose to save the group with two dying people, rather than flip a coin between that group and the group with one dying person.

Now, a claim might be disabled (or silenced) in one of two ways. It might be disabled by the strength of a competing claim to be saved, or it might be disabled by the strength of a competing claim to an equal opportunity to be saved. Given that the claim to be saved is stronger than a claim to an equal opportunity to be saved, it would be appropriate if a claimant's claim to be saved had stronger disabling power than a claimant's claim to an equal opportunity to be saved. Thus, a claimant's claim to be saved disables stronger competing claims than a claimant's claim to an equal opportunity to be saved does.

Let us consider this idea in a little more detail. Someone who is dying has a very strong claim to be saved. The strength of their claim is such that it disables lost hand and sore throat claims; it would be disrespectful to attend to sore throats in light of the dying person's claim to be saved. Thus, lost hand and sore throat claims cannot be aggregated against the dying person's claim to be saved.

However, if the dying person's claim to be saved is matched by another dying person's claim to be saved, the dying person no longer has an unmatched claim to be saved; their claim to be saved has been taken out of consideration. Thus, their claim to be saved can no longer disable the lost hand and sore throat claims. Hence, it is no longer disrespectful to consider attending to the smaller claims because they are no longer in competition with the dying person's claim to be saved.

Nonetheless, the dying person still has a claim to an equal opportunity to be saved, even though this claim is much weaker than his initial claim to be saved. His claim to an equal opportunity to be saved can still disable the sore throat claims, but it cannot disable the lost hand claims. It is disrespectful to the dying person to consider attending to sore throat claims in light of his claim to an equal

opportunity to be saved. However, it is not disrespectful to consider attending to lost hand claims in light of his claim to an equal opportunity to be saved.³⁸

We can see then how this distinction might explain the distinction between Anchoring and Competing Relevance and why they might extend different distances. Competing Relevance is based around what a claimant's claim to be saved can disable, and Anchoring Relevance is based around what a claimant's claim to an equal opportunity to be saved can disable. In this way, we might say a claim has Anchoring Relevance if it is strong enough to respectfully outweigh the strongest claimant's claim to an equal opportunity to be saved. And we might say a claim has Competing Relevance if it is strong enough to respectfully compete with another claimant's claim to be saved.

So, I have identified two potential ways in which we might begin to understand the deeper justification behind the differing notions of relevance. There is certainly a debate still to be had about the best rationale for this distinction, and which wider account of relevance gives the best explanation. However, as I make clear in the introduction to this thesis, I do not want to commit myself to any particular justification as doing so will significantly limit the conclusions of my arguments. Nevertheless, what I have shown here is that the Hybrid Approach is not ad hoc but points to a genuine moral difference and makes sufficient reference to the respect-based reasons that all relevance theorists agree on. By couching the Hybrid View in these terms, I hope the generality of my conclusions to all relevance views should be evident. I leave it to the reader to flesh the distinction out further with their own preferred rationale.

Conclusion

In summary, I have identified a fatal trilemma for both Global and Local Relevance views, showing why no such account can capture all our limited aggregation intuitions. Neither view can capture the

³⁸ In Voorhoeve's terms, the permissible personal perspective will not allow a claimant to insist upon their own rescue over the rescue of another's significantly stronger claim. However, the permissible personal perspective might allow them to prefer their own rescue over an equal opportunity of rescue for another's significantly stronger claim.

Conclusion

sore throat intuition, without abandoning the intuition that some claims might be too weak to compete but strong enough to tiebreak.

Departing from either view, I then developed a Hybrid Approach by which we can avoid this trilemma and capture all our intuitions. I distinguished between two important roles for relevance, Matching and Anchoring, and demonstrated that we need two separate relevance concepts to play both roles if we are to capture all our intuitions. I then showed that Global Relevance is well-suited to the Anchoring Role, whilst Local Relevance is suited to the Matching role, and that combining these views allows us to capture all our intuitions.

Subsequently, I applied the Hybrid Approach to a more tightly specified version of limited aggregation: *Sequential Claims Matching by Closeness*. The amended version of *Sequential Claims Matching by Closeness* has a further attractive feature in that it allows for a graded drop off in relevance, as claims get weaker. Strong claims have competing relevance with the strongest claim. Slightly weaker claims do not have competing relevance with the strongest claim but do have anchoring relevance. Even weaker claims have neither competing nor anchoring relevance with the strongest claim, but they may still be decisive in rare cases where they have competing relevance to other weak, yet relevant, claims (and can thereby counterbalance them). And lastly, we have the weakest claims, which are never relevant.

Lastly, I briefly addressed the deeper rationale behind the Hybrid Approach. I identified two possible ways of explaining the distinction. The first drew a distinction between respect-based reasons derived from actions and those derived from procedures. The second drew a distinction between respect-based reasons derived from claims to be saved and claims to an equal opportunity to be saved.

Nevertheless, one aspect that we have not given attention to in this chapter is whether Anchoring Relevance ought to be sensitive to the number of claims, and if so, how sensitivity to quantity could be incorporated into the view. For instance, we might think that a single lost finger claim is

Chapter 4 – Tiebreaks and Two Types of Relevance

insufficiently strong to decide between balanced death claims, but 100 lost finger claims do have sufficient weight to tiebreak between death claims. The view I have developed in this chapter is insensitive to numbers, and only considers the strength of a tiebreaking claim. Thus, if 100 lost finger claims are relevant and can tiebreak between death claims, then a single lost finger claim can tiebreak too. I turn to this issue in the next chapter, in which I will develop a way in which we can capture sensitivity to numbers in tiebreaking cases. We will then combine the hybrid view developed in this chapter with the findings from the next chapter, into a more comprehensive limited aggregation account in Chapter 6.

Chapter 5 – Two Approaches to Tiebreaking

Introduction

In the last chapter we saw how relevance views can be amended to capture our intuition in tiebreaking cases. By distinguishing between anchoring and competing relevance we can capture the intuition that a sore throat ought not to be decisive when deciding which of two people to save. Importantly, the hybrid approach can capture this intuition without losing our intuitions in several other important cases.

In this chapter, and the one that follows, I want to leave behind the competing cases and focus entirely on these tiebreaking cases. Whilst the hybrid view outlined in the previous chapter makes progress on previous relevance views, the account still struggles to capture a number of important tiebreaking intuitions.

This chapter will focus on some of the intuitions about tiebreaking that the relevance view cannot currently capture. Namely, we will look at cases where the number of tied, or tiebreaking, claims differ in quantity. According to the relevance view so far outlined, tiebreaking is entirely insensitive to the number of tied claims and the number of tiebreaker claims, yet we often think that whether we hold a lottery or choose one group outright ought to differ depending on the number of claims.

In response we will look at John Broome's account of fairness, and how his account can capture the sensitivity to numbers that seems appropriate in tiebreaking cases. Nevertheless, I will show that this rival account of aggregation, which I shall call the outweighing approach, has the opposite problem to the relevance approach. Broome's view is overly sensitive to the number of tied or tiebreaker claims when it ought not to be.

Seeing the two accounts juxtaposed will help us to understand why both accounts fail to capture our intuitions in these tiebreaking cases. This will lay the groundwork for a much more comprehensive account of tiebreaking that I will develop in the next chapter. This account will combine elements from the relevance and outweighing approaches to create an approach that is appropriately sensitive to the number of tied or tiebreaker claims and allow us to capture all our intuitions.

The Relevance Approach

To see in detail how relevance approaches and outweighing approaches differ, I think it would be helpful to briefly look at the basic principles behind relevance views and where these views agree. Doing so should help us to understand the core claims behind relevance views and will help me develop the view I do in the next chapter. To do so let us look at the common features of Kamm's original relevance view, and Voorhoeve's more recent formulation.

Firstly, it should be reiterated that questions about aggregation are essentially questions about the relative satisfaction of claims or the distribution of goods amongst people. Of course, such relational questions are in essence questions about fairness. Whilst not all relevance theorists talk about their views as concerning fairness, I think it is a helpful framework to understand relevance views. In particular, understanding relevance views as regarding fairness will help us see how relevance views contrast with outweighing approaches, which are explicitly about fairness.

With this in view, I think we can identify two core claims, which I put in terms of fairness, that underly relevance views:

1. Fairness is about side constraints; it rules out certain considerations. (Kamm, 2002, p. 686) (Voorhoeve, 2014, p. 70)
2. Fairness involves giving "*significant priority* to those who are the worst off" (Kamm, 2002, pp. 686-687). (Voorhoeve, 2014, pp. 72-75)

The Relevance Approach

It is important that we understand these claims in tandem here. Side constraints put absolute limits on what we are allowed to take into consideration, or what we are allowed to do. A common example of a side constraint would be an inviolable right, like the right to life (though not all side constraints are rights). According to this view, side constraints cannot be outweighed.³⁹

Now for relevance theorists, priority involves side constraints. Rather than say that the worst off should be given extra weight, we say that the concerns of the worst off provide side constraints and rule out the consideration of much weaker concerns. In the Sore Throat case, the priority owed to the dying people puts a side constraint on what we can take into consideration. Thus, the sore throat becomes an irrelevant utility, is dismissed from consideration, and so we ought to flip a coin between the two groups.

Now, as we have already seen, there are several different explanations for when and why a stronger claim has sufficient priority over a weaker claim such that side constraints are generated and the weaker claim becomes irrelevant. Depending on the account, we might say that it occurs when it would be disrespectful to the stronger claimant to take the weaker claim into consideration (Kamm, 2007, p. Chapter 10); or when it would not be permissible for the weaker claimant to prefer their own claim to be saved over the stronger claimant's (Scanlon, 1998, p. Chapter 5); or when we lose sympathy for the weaker claim in light of the death claim (Voorhoeve, 2014) (2017).

As per previous chapters, because I want to show that my conclusions apply to all relevance views, I will not engage in the debate over which of these explanations is best. Nevertheless, it should be noted that they all reference the importance of considering each person's personal perspective. The exact detail of how to account for the personal perspective does not matter here, so I will set this aside; we only need the loose detail. In essence then, we consider the partial concern of each individual, what is at stake for them, and then decide which claims have sufficient strength to form

³⁹ Though they can be overridden by other considerations, i.e., other side constraints.

part of a respectful justification to each of those individuals whom we do not save. In this way, relevance is determined not by the number or quantity of claims, but only by the strength or quality of the claims.

For instance, in Kamm's Sore Throat case, when we consider the perspective of the death claimant in Group A, we recognise that a sore throat could not form part of a respectful or acceptable justification not to give them a chance to be saved. Nor is it permissible for someone to prefer their own sore throat to be cured over someone else's chance of having their life saved. Nor can we sympathise with a sore throat in the context of a life. Thus, sore throat claims lose their duty grounding force and are made irrelevant by the presence of the death claim.

Large Scale Cases

We should be familiar enough by now with the way relevance views work and the justification behind them. So, there is no need to go into more detail. As such, let us look immediately at some of the issues for the relevance approach. Firstly, let us look at a case that Kamm, perhaps unwittingly, identifies for relevance views:

"Now suppose 1000 people are on one island and 1001 people are on another. Here, I believe, it may even be correct to ignore the difference of one life. If so, then in this context the one life has become an irrelevant utility. This would be true even if, in a choice between two on one island and three another, the additional life should be determinative of whom we save.

[...] This is because we understand that the equal 50 percent chance to survive is important to the [1000], and that from their points of view it is not just as good if the other [1000] people survive. We think that preserving this equal chance of so many is more important than saving one extra person outright." (Kamm, 1998, p. 103)

For consistency, let us put this into the same terms as previous cases:

The Relevance Approach

Kamm's Large Scale Case: You can save Group A or Group B. In Group A you have 1000 people who will die. In Group B you have 1001 people who will die.

This case puts serious pressure on the relevance account. Relevance views have a dilemma in this case.

Relevance theorists have two options:

1. Say that the 1001st life is relevant to our decision and decide in favour of Group B.
2. Say that the 1001st life is irrelevant and give the two groups equal chances to be saved.

Neither option here is particularly appealing. The first option says the right thing about the 1001st life; it seems clear that if the other lives matter, then the 1001st life also matters. However, as Kamm highlights, our intuition is that if the numbers of balanced death claims are sufficiently large then we ought not to let the one additional death claim be decisive. The extra death claim in Group B is insufficient to be decisive in favour of that group. Relevance theorists could thus bite the bullet in these large-scale cases and say that no matter how many claims we have on each side, a single death claim can be decisive.

Now, biting the bullet is not an illegitimate strategy but any view that can capture our intuitions is *prima facie* stronger, and ought to be given greater attention. This is especially so when the whole field of limited aggregation is so driven by our intuitions around these important cases. As such, I shall set the first of these two options aside.

Now the second option is the option that Kamm seems to indicate she would take:

“[T]he number of others whose lives are at stake may make relevant or irrelevant an additional utility that is irrelevant or relevant in other contexts.” (Kamm, 1998, p. 103)

So, according to Kamm, quantity might factor into questions of relevance. The strength of a claim might vary against the background context, and in this context the strength of the extra claim is

completely irrelevant. In the Large-Scale case, the extra life is made irrelevant because we ought to prioritise giving the greater number an equal chance of being saved.

But this is clearly not right. It is just wrong that a death claim could be made irrelevant by the presence of other death claims. The 1001st death claim is just as important as the other death claims. To treat the 1001st claim as irrelevant violates *Equal Consideration for Equal Claims* - i.e., the principle that “all claims of equal strength ought to be given equal weight in determining which group to save” (Tomlin, *On Limited Aggregation.*, 2017, p. 243). If a claim of a certain strength counts in that context, then all claims of that strength count in that context. In other words, it would be wrong to consider some death claims as mattering and others as not.

This is made especially problematic if we can identify the claim that is irrelevant. For instance, suppose we have two equally matched islands with 1000 people on each, and so we decide to run a lottery to decide whom we save. But just before we do, pirates maroon an extra person on the second island. To dismiss this extra person’s claim as irrelevant, as not counting, is incredibly disrespectful. Kamm and other relevance theorists could bite the bullet here, but it strikes me that this is a bigger bullet to bite than to say the 1001st claim is decisive in this case. More importantly, if a view can capture this intuition without making the 1001st claim irrelevant, that will be a significantly stronger view.

Furthermore, even if we do decide to bite the bullet in this case, the relevance view would still need significant altering. On a theoretical level, relevance views have been developed in the literature to consider only the claims of individuals, not of collectives, when determining relevance. Remember that Kamm, Scanlon and Voorhoeve all argue (in slightly different ways) that relevance is determined by comparing how much the stronger and weaker claimants have to lose. This is because relevance is intended as an anti-aggregative notion; one that uses respect for the separateness of persons, the personal perspective and/or partial concern to determine whether a claim counts.

The Relevance Approach

But in the Large-Scale case, there is no individual who could complain about the inclusion of the 1001st death claim. Nor does the 1001st claimant have a duty to set aside her claim to be saved; it would be perfectly permissible for her to press her claim and choose to save her own group than to leave which group is saved up to chance. Thus, if relevance theorists want to dismiss this claim as irrelevant, they are going to need not just to just provide a good justification for doing so, but also to find a whole new approach to understanding relevance. Appeals to the separateness of persons, the personal perspective and/or partial concern of individuals will be insufficient to explain what makes a claim relevant.

Of course, this is not to say that a relevance theorist could not take this approach. However, anybody who does take this approach is going to have to abandon many of the foundational principles and much of the progress made by relevance theorists. Such a project would be very large, but perhaps some might think the project worth the effort. If so, they are welcome to take it up. As for me, I believe such a project would be to throw the baby out with the bathwater, and a serious overreaction to a difficult case. It is highly unintuitive to consider the 1001st life irrelevant in the first place.

Many Tiebreakers

Now let us look at another case that makes a very similar point; but, instead of a large number of tied claims, we have a large number of tiebreakers:

Many Tiebreakers (1): You can save Group A or Group B. In Group A you have a death claim and in Group B you have a death claim and N lost finger claims.

Who should we save in this case? My strong intuition, which I imagine many will share, is that when N is small, we should give equal chances to both groups, but that when N is sufficiently large we should save Group B. It seems to me that a single lost finger claim is not strong enough to tip the

balance in favour of Group B. However, a sufficiently large number of lost finger claims would tip the balance in favour of Group B.

Now one might disagree with me about the specific claim I use here. Perhaps, one thinks that a single lost finger claim is sufficiently strong to be decisive, or that no number of lost finger claims could be decisive. If so, the reader should consider different strength claims. These might be stronger lost hand claims; alternatively, they might be weaker broken finger claims. My point is simply that there are going to be some claims too weak to tiebreak by themselves, but that could be decisive in concert with a large number of other equally weak claims.

If there is anything to this intuition (and I think there is), then this could be quite problematic to the relevance account for tiebreakers. Relevance views have to give the same answer, no matter how large N is; the lost finger claims are either relevant or irrelevant. If their claims are irrelevant, then no number of claims will be able to tiebreak. If lost finger claims are relevant, then even a single one would tip the balance in favour of Group B.

Now the relevance theorist might try to resist or explain away the intuition in this case. She might argue that a claim that is strong enough to aggregate and outweigh a death claim should be able, even in small numbers, to tiebreak between death claims. Once we develop the relevance theory in this way, it might seem like we have good reasons to think that relevant claims even in small numbers can tiebreak.

There are two things to say here. Firstly, this response does not actually explain away the intuition. It might give us theoretical reason to abandon the intuition, but we would still need to explain why we have the intuition. This might be easier if this is the only such case of this intuition, but *Kamm's Large Scale Case* also demonstrates that we think a single relevant claim might not always be enough to tiebreak. It is this very reasoning that these cases are supposed to apply pressure to. As before, biting the bullet here is not an illegitimate strategy and we might think that the relevance

The Relevance Approach

theorist's reasoning here might also lessen the bullet's sting. However, the number of bullets is adding up, and a view that could avoid these bullets entirely would still be preferable.

Secondly, and more importantly, the relevance theorist's reasoning here is flawed anyway. It assumes that the only claims that could be relevant in a tiebreaking situation are claims that are also relevant in an outweighing situation. I.e., that the only claims that could tiebreak between death claims are claims that in sufficient numbers could outweigh a death claim. But as we saw in the last chapter, relevance in the tiebreaking and outweighing cases come apart, and should be treated differently.

To further emphasise why we should think quantity matters, consider the weakest claim that could tie break between two death claims by itself. Suppose this is a lost hand claim. Now consider a claim slightly weaker than a lost hand claim - say a claim against losing four fingers. On this view, the four-finger claim could not tiebreak between death claims - no matter how many people had such claims. This is quite a strange conclusion. A more natural understanding would be to say that a single lost hand claim could tiebreak between death claims, but several four-finger claims are needed to tiebreak. This would capture a more gradual weakening of tiebreakers, rather than a sharp line delineating relevance.

The only option that avoids biting the bullet in *Many Tiebreakers (1)* is to argue that the lost finger claims are irrelevant when N is small, but that they are relevant when N is large. But I just do not see how such a view would work. The death claimant's claim is not stronger when N is small, nor are the lost finger claims individually stronger when N is large. If a claim's being irrelevant is based on partial concern, or what is fair to individual claimants, then it can only be based on the quality of the claims not on quantity. As we have seen, relevance views are predicated on the very notion that only quality matters when deciding whether a claim is relevant or not. If a claim's relevance could be determined by the quantity of claims it either competes with or competes alongside, then we open the door to the tiniest of claims in large enough numbers swamping the decision - an outcome

relevance views were designed to prevent. Thus, those who want to maintain a relevance approach seem wholly unable to capture our intuition in this case.

The core problem for relevance theorists is that they only have one way to capture tiebreaking cases. If we want to give two unequal groups equal chances, then we have to say that the difference between the two group is irrelevant. This also means that relevance views are entirely insensitive to the number of claims. In this way, relevance views do not provide us with the conceptual tools needed to be able to explain the Large-Scale and Many Tiebreakers cases in the right way.

The Outweighing Approach

Let us now turn to Broome's view. It is important to do so for a few reasons. Firstly, Broome's account is immensely impressive and highly influential. Thus, to not consider it would be to ignore an important element of the literature. Secondly, considering Broome's outweighing approach against the relevance approach will help us to understand the issues with both, and will allow us to develop a comprehensive hybrid view in the next chapter.

Now, whilst Broome is a pure-aggregationist, he agrees with Kamm's conclusions in both the Sore Throat and Large-Scale case. However, he disagrees with the way she gets to her answers. As a pure-aggregationist he disagrees with her principle of irrelevant utilities:

“Kamm argues that some goods in some circumstances count for nothing. But I believe that all goods always count for something.” (Broome, 2002, p. 739)

Instead, he offers an outweighing approach to fairness which is characterised by these two core claims:

1. Fairness is an end state, and only provides pro tanto reasons. (Broome, 1998, p. 2)
2. Fairness requires “that *claims should be satisfied in proportion to their strength.*” (Broome, 1990, p. 95).

The Outweighing Approach

We can see here that Broome's approach to fairness is at complete odds with the relevance approach outlined in the first section. Instead of side constraints, fairness is an end state value. Instead of being concerned with priority, fairness is concerned only with proportionality.

Now, I should make clear the role of claims in Broome's theory. Whilst I set out in the introduction to the thesis that I primarily talk of claims for reasons of simplicity and that the view I develop is not reliant on claims talk, Broome's view does rely on claims. For Broome fairness is only concerned with claims, and not wider reasons. If John might benefit from me giving him £10, that gives me a reason to give it to him. But if John has no claim to my £10 whilst I owe Frances £10, then an equal distribution would be unfair. That Broome's account of fairness involves only claims is contentious (Tomlin, 2012), but we can set such concerns aside as our focus is on the proportionality and, more importantly, the outweighing aspects of Broome's view.⁴⁰ For simplicity, I shall also set aside the conditions that might generate a claim and change the strength of a claim (for instance, having worked hard for the resource). Instead, I shall assume in all my examples that the people have claims to the good, and that their claims' strengths are entirely based on their needs.

Thus, fairness for Broome involves the proportional satisfaction of claims, such that, if your claim on some divisible good is twice as strong as mine, then you should get twice as much as I do. So, fairness involves a proportional distribution of resources where there are rival claims on these resources, or, in some circumstances where the resource is indivisible, a proportional chance of receiving the resource. Thus, even very weak claims remain relevant and fairness requires their proportional satisfaction.

Importantly, proportionality is not a side constraint that puts limits on what we are allowed to do, nor do considerations of proportionality ever make other considerations irrelevant. Instead, a

⁴⁰ More broadly my argument against outweighing approaches to fairness does not rely on Broome's use of claims. We could reconstruct the proportionality and outweighing aspect of his account without reference to claims and the account would still fail to capture our intuitions in the cases I shall outline below. In this way, we can engage with Broome's view on its own terms and it will still fail.

proportional satisfaction of claims is what fairness requires; it has value in and of itself. However, it is not something we should aim for at the cost of every other consideration. Thus, the reasons given by fairness must be weighed against other reasons and this means that fairness concerns can be outweighed if the other concerns are strong enough.

Of course, to be able to weigh the fairness concerns against the other concerns we need to know how to determine the total value of the fairness concerns. To do so we need to understand the hierarchical nature of Broome's view. For Broome there are two types of fairness, outcome fairness and procedural (or surrogate) fairness and these types of fairness will sometimes have to be weighed against each other:

The likelihood of [a] less fair result will have to be weighed against the contribution to fairness of the lottery [i.e., the fair procedure] itself. (Broome, 1990, p. 99)

Now, outcome fairness is achieved when each claim is satisfied in proportion to its strength. For instance, if I owe both John and Frances £10, but I only have £10, then the fair thing to do is to give them both £5. Whilst neither have got what they are owed, they have their equal claims satisfied equally and thus there is no unfairness between them.

However, now consider that I only have a single indivisible £10 note. Thus, I cannot give John and Frances £5 each. Outcome fairness would only be achieved if I gave neither of them the £10. This is the only way to treat them equally. However, the outcome unfairness of giving either John or Frances £10 and the other nothing is outweighed by the need to satisfy one of their claims. Thus, it is better here to be unfair and satisfy one claim, than to be fair and satisfy no claims. This is what it is for fairness to provide a pro tanto reason.

Nevertheless, in this indivisible case, we can mitigate the outcome unfairness by introducing procedural fairness. According to Broome, when a good is indivisible, procedural fairness requires that we give each claimant a proportional chance of getting that good. For instance, John and

The Outweighing Approach

Frances each have a claim to an equal chance of receiving the £10 note. If I decide to give John the £10 because I like him more, Frances has a complaint that she has been treated (procedurally) unfairly as she did not have a fair chance of receiving the £10 note. The procedurally fair thing to do then is to run a lottery – or in this case flip a coin.

Now, it is important to understand that on the outweighing approach the badness of unfairness is attributed to the claimants:

Unfairness, as I have described it, is plainly an individual harm. There is unfairness if someone's claim is satisfied less than in proportion to its strength. Since a claim is a duty owed particularly to the person, the unfairness is plainly suffered by that person. If, say, people have equal claims to the satisfaction of needs, and some people have their needs less well satisfied than others do, then those people are suffering unfairness. (Broome, 1991, p. 198)

Thus, if I give John the £10 without flipping a coin, the badness of the outcome and procedural unfairness is suffered by Frances. She is the victim of the unfairness and has a complaint that she was not given a fair amount of the good nor a fair chance of receiving the good.

Frances and John thus have three different types of claims. Firstly, they each have a claim to the good - in this case, £10. Secondly, they each have a claim to a fair proportion of the good – in this case, £5 each. Thirdly, when they cannot receive their fair proportion of the good, they each have a claim to a fair chance of the good – in this case, a 50% chance of getting the £10. In this way, procedural fairness is derived from outcome unfairness (hence why Broome calls it a surrogate).

These two latter claims, we call 'fairness-claims'. As Broome notes, someone might have a fairness-claim even if they have no claim to the good itself (Broome, 1990, p. 97). A child is not owed sweets from his parents, but if his parents are giving out sweets he has a fairness-claim to an equal number

of sweets as his sister.⁴¹ Similarly, if there are an odd number of sweets, the boy has a fairness-claim to an equal chance of the last sweet.

Fairness-claims are a vital and underappreciated aspect of Broome's account. Saying that each individual has a fairness-claim allows Broome's to locate the badness of unfairness in each individual who does not get their proportion of the good or their proportional chance of the good. Unfairness then is not just a bad state of affairs, but a badness borne by those who are treated unfairly. This means that the badness of unfairness increases as more people are treated unfairly. In turn, this allows Broome to capture the Large-Scale cases as we will see later.

To summarise, the four options in the case where I can only give John or Frances £10 are:

Give Nothing: No utility, No outcome unfairness, No procedural unfairness

Give to John: Maximum possible utility, Outcome unfairness for Frances, Procedural unfairness for Frances

Give to Frances: Maximum possible utility, Outcome unfairness for John, Procedural unfairness for John

Flip a Coin: Maximum possible utility, Outcome unfairness for either John or Frances (but not both), No procedural unfairness

In this case, then we can see why we ought to flip a coin. The utility of distributing the £10 is greater than the outcome unfairness it generates. The outcome unfairness of each of the latter three options is the same (one person suffers the badness of not having their claim to a proportional amount of the good satisfied), but *Flip a Coin* also has no procedural unfairness, so we ought to choose that option. It is important to note that in this case the procedural and outcome unfairness

⁴¹ Patrick Tomlin rightly points out that this means that fairness generates claims, and that claims cannot thus be the only subject of fairness (Tomlin, 2012, pp. 206-207). This puts further pressure on Broome's use of claims, but it does not put pressure on the proportionality or outweighing aspect of Broome's view which is our focus here.

The Outweighing Approach

do not come into conflict but in other cases they will, and we will have to weigh procedural fairness against outcome fairness.

With Broome's account now in view we can return to the Sore Throat, Large-Scale and Many Tiebreaker cases. First, let us look at the sore throat case. The three options (from now on we will ignore the *Give Nothing* option) we have are:

Give to Group A: 1 life saved, Outcome unfairness for 1 life and 1 sore throat, Procedural unfairness for 1 life and 1 sore throat

Give to Group B: 1 life and 1 sore throat saved, Outcome unfairness for 1 life, Procedural unfairness for 1 life

Flip a Coin: 1 life (and $\frac{1}{2}$ a sore throat)⁴² saved, Outcome unfairness for 1 life (and $\frac{1}{2}$ a sore throat), No procedural unfairness

Give to Group B dominates (i.e., is stronger in all elements) *Give to Group A*. So, we take that option out of consideration. To decide between *Give to Group B* and *Flip a Coin*, we then simply compare the utility, outcome fairness, and procedural fairness. We choose *Flip a Coin* if the procedural unfairness outweighs the extra utility and reduced outcome unfairness in *Give to Group B*. In this case, the procedural unfairness suffered by the dying person does outweigh the extra utility and outcome fairness to the sore throat claimant, and thus we *Flip a Coin*.

Now let us look at the options in the Large-Scale case. Remember that because Broome identifies unfairness as an individual bad this means that the aggregate "badness of unfairness is determined

⁴² We do not know exactly the consequence if we flip a coin as there is a half chance that the sore throat will be saved anyway. For simplicity, to capture this uncertainty I have used expected utility style calculation. The expected utility and expected unfairness of flipping a coin is $\frac{1}{2}$ a sore throat saved and $\frac{1}{2}$ the outcome unfairness to a sore throat, respectively. There might be other options, for instance if one is heavily risk averse, we could instead separate *Flip a Coin* into its two states and compare the worse state with *Give to Group B*. Either way, if the procedural unfairness is strong enough, it will outweigh even the worst state of *Flip a Coin*.

by the number of persons who are the victims of unfairness.” (Hirose, 2015, p. 192). Thus, in this case we have the following options:

Give to Group A: 1000 lives saved, Outcome unfairness for 1001 lives, Procedural unfairness for 1001 lives

Give to Group B: 1001 lives saved, Outcome unfairness for 1000 lives, Procedural unfairness for 1000 lives

Flip a Coin: 1000 (and a $\frac{1}{2}$) lives saved, Outcome unfairness for 1000 (and a $\frac{1}{2}$) lives, No procedural unfairness

Again, *Give to Group B* dominates *Give to Group A* so we dismiss that option. In this case, we choose *Flip a Coin* if the procedural unfairness to the 1000 is greater than $\frac{1}{2}$ a life saved and $\frac{1}{2}$ the outcome unfairness to one life. The outweighing approach says it does, and thus we ought to flip a coin in the Large-Scale case.

Importantly, this means Broome’s view can capture Kamm’s intuition in the Large-Scale case without making any claims irrelevant. The extra life is relevant and taken into consideration, however the extra life claim is not competing with the procedural fairness owed to one person, but to the aggregated fairness-claims of a thousand other people. These fairness-claims together outweigh the extra good that would be done by directly saving the extra death claim.

This also means that the outweighing approach can capture the Large-Scale case whilst explaining that in a case between saving one life or saving two, we should always save the two. In the smaller case, the single fairness-claim in the first group does not outweigh the extra claim to be saved in the second. Unlike the relevance view, this also means the outweighing approach provides a good explanation for why we should run a lottery in the Large-Scale case. Broome has a way to legitimately appeal to the numbers in the Large-Scale case in a way that the relevance view is unable to.

The Outweighing Approach

Lastly, let us look at the Many Tiebreakers case:

Give to Group A: 1 life saved, Outcome unfairness for 1 life and N lost fingers, Procedural unfairness for 1 life and N lost fingers

Give to Group B: 1 life and N lost fingers saved, Outcome unfairness for 1 life, Procedural unfairness for 1 life

Flip a Coin: 1 life and $(\frac{1}{2})N$ lost fingers saved, Outcome unfairness for 1 life and $(\frac{1}{2})N$ lost fingers, No procedural unfairness

Again, *Give to Group B* dominates *Give to Group A*, no matter how large N is. However, we can see that whether we ought to *Give to Group B* or *Flip a Coin* comes down to the size of N . If N is small, then the procedural unfairness outweighs both the lost utility and outcome unfairness to the lost fingers and so we ought to *Flip a Coin*. If N is large, the lost fingers outweigh the procedural unfairness and so we ought to *Give to Group B*. Broome's view can thus capture our intuition in this case too: his view is appropriately sensitive to the numbers.

We might think then that we ought to simply adopt an outweighing approach and be done with it. However, the outweighing approach suffers from cases that are (at least) as problematic as that which afflict the relevance approach.

Strong Tiebreakers

Before I offer my own critique, I want to briefly explore Brad Hooker's as it will help us see how my critique works:

[S]uppose that there is available some limited quantity of medicine and that this medicine cannot be divided without rendering it ineffective. Suppose your claim on the medicine comes from the fact that you need it to save your life, and my claim on it comes from the fact that I need it to save my little finger. Suppose an average life is something like a thousand times more important than a little finger. So should the matter of who gets the

medicine be decided by a lottery in which you have a 999/1000 chance of winning and I have a 1/1000 chance? Given that your claim is so much stronger than mine, how could it be right to take any risk that I rather than you might end up with the good?. (Hooker, 2005, p. 349)

Hooker argues that, on Broome's view, fairness requires we hold a lottery here, but that this seems unfair. However, I think Hooker's characterisation of Broome's view here misses the distinction Broome makes between fair results and the fairness a lottery provides. Hooker seems to approach Broome as having a divided view of fairness: that when a good is divisible fairness requires a proportionate satisfaction of claims, and when the good is indivisible fairness requires a lottery. But if this were the correct interpretation, we would not have to weigh fair results against the fairness the lottery provides, as Broome says we ought to. (Broome, 1990, p. 99)

Thus, as Christian Piller shows (2017, pp. 229-232), Broome is not committed to this consequence. In Hooker's case we must weigh the procedural fairness of the lottery against the risk of serious outcome unfairness to the death claimant. We must ask "[w]hat is more unfair: reducing a small chance to 0 or risking that a much weaker claim is satisfied?" (Piller, 2017, p. 231). Broome would answer that it is more unfair to risk not saving the life, than reduce the chance of a single lost finger claim being saved to 0.⁴³

Inspired by Hooker's critique, let us now consider a much more devastating problem which Broome's view cannot avoid. This is a case that both relevance theorists and utilitarians will agree on, but that the outweighing approach gets very wrong.

Strong Tiebreakers Case: You can save Group A or Group B. In Group A you have N number of sore throat claims. In Group B you have N number of sore throat claims and a death claim.

⁴³ Whilst Broome can avoid saying we ought to have a lottery in Hooker's case, he does still have to say that there is a pro tanto reason to hold a lottery. In so far as Hooker's case shows that there is no unfairness here at all - that there is no pro tanto reason to hold a lottery whatsoever, even if that reason is outweighed by other considerations - Hooker's argument still works. However, whilst this is correct it is probably a bullet Broome would be willing to bite.

The Outweighing Approach

In this case we should definitely save Group B. There is somebody dying in Group B and there are an equal number of sore throats in both groups. It would be beyond the pale to run a lottery in this situation. Pure-aggregationists, anti-aggregationists and limited aggregationists of all shapes and sizes will all agree that we ought to save Group B. The death claim, if not strong enough to win outright by itself, must, at least, tiebreak in favour of Group B.

Yet Broome's outweighing approach to fairness will be forced to run a lottery if the number of sore throats is large enough. We could even increase the number of death claims in Group B to 1000, and there would still be some number of sore throats such that the outweighing approach to fairness would require a lottery. A truly absurd consequence.

To see why we need to consider our options in the Strong Tiebreaking case. For our purposes let us suppose that the average life is 1000 times more important than a sore throat (in reality it is magnitudes more important, but we should keep numbers manageable here). Given that Broome argues against pooling chances (Broome, 1998, pp. 4-5), a fair lottery between the two groups would give Group A a 1/1000 chance of being saved and Group B a 999/1000 chance of being saved.⁴⁴ Thus, the three options we have are:

Give to Group A: N sore throats saved, Outcome unfairness for N sore throats and a life,

Procedural unfairness for N sore throats and a life

Give to Group B: N sore throats and a life saved, Outcome unfairness for N sore throats,

Procedural unfairness for N sore throats

⁴⁴ It is not actually entirely clear what Broome's view would say a fair lottery requires in heterogenous cases like this as Broome does not address them. Broome's view is essentially comparative, and it is unclear what proportionality requires when there are different strength claims in the same group. I choose here to interpret Broome as deciding the weights of the lottery based on the strongest claims in each group as this seems to fit best with his argument against pooling chances. However, this approach has other flaws, as do the other ways we might determine what proportionality requires in heterogenous cases. This is an issue I expect to address elsewhere, but whichever account we take will not affect my argument here. My argument is against lotteries in these cases full stop, and does not rely on how the lottery's weighting is determined.

Run a Lottery N sore throats (and 999/1000 of a life) saved, Outcome unfairness for N sore throats (and 1/1000 of a life), No procedural unfairness

Of course, *Give to Group B* dominates *Give to Group A*, so we ignore this option. To choose between *Give to Group B* and *Run a Lottery* we then compare the differences between the two groups. If we *Give to Group B*, we save 1/1000th of an extra life and avoid the outcome unfairness to 1/1000th of a life compared to *Run a Lottery*. But we also have procedural unfairness to N sore throats. Thus, we *Give to Group B*, if and only if saving 1/1000th of a life and avoiding outcome unfairness for 1/1000th of a life, outweighs the procedural unfairness to N sore throats. To put it in Piller's terms, we must ask what is more unfair: reducing many small chances to 0 or risking that the strongest claim is not satisfied?

The answer on Broome's view will depend on how many small chances there are. The unfairness of the strongest claim not being satisfied remains the same as N increases. Whereas the badness of unfairness of many small chances being reduced to 0 increases as N increases. Thus, on Broome's outweighing approach, if N is large enough, we ought to run a lottery in the Strong Tiebreakers case. This is clearly an absurd consequence.

Broome cannot avoid the consequence here in the way his view can in Hooker's case. In Hooker's case, Broome can argue the outcome unfairness outweighs the procedural unfairness. But in the Strong Tiebreakers case the outcome unfairness has already been factored in. As N increases, the outcome unfairness is bounded, whilst the procedural unfairness increases unbounded.

To make the problem even clearer, consider how the Large-Scale case and the Strong Tiebreakers case have the same structure. In both cases there are a large number of equal claims on both sides, all of which have procedural fairness-claims (i.e., claims to a chance of the good). Thus, without the tiebreaking claims, in both cases, we ought to run a lottery. The outweighing approach says in the Large-Scale case that these aggregated procedural fairness-claims will outweigh the extra life and so

The Outweighing Approach

we ought to flip a coin. But without the ability to say that the sore throats are irrelevant, this approach will also tell us to run a lottery in the Strong Tiebreakers case.

In fact, we could start with the Large-Scale case, and slowly replace the balanced death claims one by one, with an increasingly larger number of increasingly weaker claims:

Table 1

	Group A	Group B
Large-Scale case	1000 death claims	1001 death claims
Large-Scale case 2	2000 paraplegia claims	2000 paraplegia claims + 1 death claim
Large-Scale case 3	4000 lost limb claims	4000 lost limb claims + 1 death claim
...
Strong Tiebreakers	N sore throat claims	N sore throat claims + 1 death claim

If we are to say that we ought to run a lottery in the Large-Scale case but not in the Strong Tiebreakers case, we need to show that at some point the weaker balanced claims are irrelevant. As we have seen, the outweighing approach does not have a mechanism for this. So Broome is stuck with running a lottery in all these cases. Given how devastating the Strong Tiebreakers case is, Broome is better off abandoning this theory of fairness completely (and biting the bullet in the Large-Scale case and in Kamm's Sore Throat case) than having to bite the bullet in the Strong Tiebreakers case.

Potential Responses

But perhaps Broome can avoid this consequence. There is some evidence that Broome only intended for his view to cover the unfairness done to the strongest claims:

"A lottery should be held when, first, it is important to be fair and, second, the candidates' claims are equal or roughly equal." (Broome, 1990, p. 99)

Broome here wants to limit the role of lotteries. If we were to understand his claims here as being deontic, as providing side constraints on when we can use lotteries, then this would be sufficient to avoid Strong Tiebreakers. However, Broome wants to consequentialise these deontic principles, so this is not what he means.

When Broome says that lotteries should only be held when it is important to be fair, he simply means to repeat what he has already explained; that “[i]n some circumstances, no doubt, it will be very important to be fair, and in others fairness may be outweighed by [the good]” (Broome, 1990, p. 96). However, as the Strong Tiebreakers case shows, the extra goodness of saving the life will not always outweigh the aggregated fairness-claims, and so this does not help him here.

Similarly, when Broome says that lotteries should only be held when claims are equal or roughly equal, he is not making an appeal to deontic or relevance conditions. Instead, he is repeating the point that lotteries generate outcome unfairness, and when claims are wildly unequal a lottery (even a weighted lottery) risks generating a large amount of outcome unfairness, and the benefit in procedural fairness does not outweigh it. (Piller, 2017, pp. 229-232) However, as the Strong Tiebreakers case shows, this outcome unfairness will not always outweigh the procedural fairness when claims are wildly unequal. The difference in outcome unfairness between *Save Group B* and *Run a Lottery* in Strong Tiebreakers remains the same when the number, N , of sore throats in each group increases. However, the difference in procedural unfairness between the two options increases linearly as N increases. Thus, at some stage the procedural unfairness will inevitably outweigh the outcome unfairness, even if the claims are wildly unequal.

The only other way that Broome might avoid the consequence in the Strong Tiebreakers case is to say that I have weighed the outcome unfairness and procedural fairness incorrectly. Now, Strong Tiebreakers only requires (A) that outcome unfairness to the dying person remains bounded as N increases and (B) that procedural unfairness to the sore throat claimants increases without bound as N increases. Thus, Broome has two options: deny (A) or (B).⁴⁵

To deny (A), Broome could say that as the number of sore throats increases, the outcome unfairness of the death claimant not being saved also increases. Thus, the fairness owed to the sore throat

⁴⁵ Broome could also say that outcome unfairness is lexically prior to procedural unfairness. But if this were the case, procedural fairness could never outweigh outcome unfairness and lotteries would only ever be appropriate between identical claims or groups of claims.

The Outweighing Approach

claimants could never outweigh the fairness owed to the death claimant. However, this is not how Broome sees outcome unfairness working. Remember that on the outweighing approach unfairness is done to each person. This is what allows the badness of unfairness to aggregate and outweigh other considerations. If Broome were to say that the badness of the outcome unfairness altered with the number of other claims, then he would lose the simple explanation in the Large-Scale case and he would need to alter his view significantly to explain how the outcome unfairness changes.

To deny (B), Broome could say that there is a limit to total procedural unfairness as N increases, either, because of diminishing marginal returns in the procedural unfairness as extra sore throat claims are added, or because there is a cap to the total badness of procedural unfairness. Again though, neither of these explanations are particularly Broomean. If procedural unfairness is suffered by each individual who does not receive a chance to be saved, then as the number of sore throats increases so must procedural unfairness. Otherwise, the latter sore throats become irrelevant. Similarly, if there were a cap on the total badness of procedural unfairness, then (if the other considerations were strong enough) lotteries would never be permitted, including in Large-Scale cases.

More importantly, in correspondence Broome confirmed that his account would suggest running a lottery in the Strong Tiebreakers case. He agrees with my analysis of the case; however, he is happy to accept this consequence. He argues that the intuition in this case is insufficient to ground an objection to the view:

“You tell me that fairness among people with sore throats could not outweigh the value of a life, however many of them they [sic] are. I’m supposed to accept that on grounds of intuition. But I don’t. This quantifier ‘however many there are’ often appears in moral philosophy, but I don’t think it is a suitable object for intuition. I see no reason to think that a lot of sore throats cannot be enough to outweigh a life. A life is not infinitely valuable, of course.” (Broome, personal correspondence, November 11th, 2022)

Broome's response is that we cannot trust our intuitions in these large number cases. I assume the reason being that human faculties struggle to cope with large numbers, and so we are prone to error. Now, I shall not enter the debate here about what is, or is not, a suitable object for intuition. Doing so would alter this chapter's focus completely. However, for those of us who put more stock in our intuitions, Broome's response is unlikely to be convincing. If the argument that we cannot trust our intuitions in large number cases collapses, then Broome's whole view of fairness will also fall like a house of cards.

Nevertheless, I do think Broome's objection has more strength to it in the Strong Tiebreakers case, because of how unrealistic the case is. Equal sore throats are never going to come into conflict with a death like this, and definitely not in very large numbers. So, it might be helpful to consider a more realistic case:

Scanlon-Style World Cup Case: You are in the transmission room for the World Cup final, where 4 billion people are watching. Half are watching via television transmission, whilst the other half are watching online. Jones gets trapped under electrical equipment and he is in terrible pain. In getting trapped he has also pulled out the wires transmitting the World Cup to everyone. We can free Jones, and in the process restore connection for everyone watching on their TV. Or we can leave Jones and restore connection for everyone watching online.

In this case, the intuition that we ought to save Jones, and restore the connection to those watching the football on the TV, is even clearer. However, Broome of course, must say that we ought to run a weighted lottery; the fairness owed to those watching online aggregates and outweighs the extra utility of helping Jones outright.

There are a couple of ways to make this case even worse for Broome. First, suppose there are only two options in this case: flip a coin or save Jones. Flipping a coin is not even procedurally fair on Jones, as he does not even have a chance to be saved in proportion to the strength of his claim.

The Outweighing Approach

However, the procedural unfairness to Jones will also be outweighed by the procedural fairness to the very large number of TV watchers. Thus, Broome's view requires us to flip a coin and in this case the lottery does not even favour Jones.⁴⁶

The second way to make this case worse is to turn it into a multiple-choice case. Suppose we can save Jones and restore connection to one person watching online or restore connection to only one of the other 4 billion people watching. Let us stipulate that Jones has a claim 1 million times stronger than each of the 4 billion watching, he then ought to have a 1 million times greater chance to be saved than each individual. But the goodness of saving him outright cannot outweigh the aggregated badness of the unfairness to the 4 billion of not having a chance to be saved. Thus, we ought to run a lottery. But this lottery will only give Jones a 1 million in 4 billion (i.e., a 1 in 4,000) chance of winning the lottery, because of the sheer number of small claims. This conclusion is clearly terrible.

I think these cases also indicate that our intuition in these cases is not disturbed by the large numbers. We simply intuit that when we can attend to equal numbers of weaker claims, that fairness between those much weaker claims should not weigh against much stronger claims. We then think that this intuition holds, even when there are large numbers. The large numbers are not central to the intuition; they simply demonstrate a case where Broome's view diverges from the relevance view and cannot capture our intuition.

Fundamentally, Broome cannot capture these cases because he tries to consequentialise fairness. Of course, the problem with all consequentialist views (as has been shown repeatedly) is that very large numbers of very minor considerations will aggregate and overwhelm much larger considerations. As these cases show, this problem persists even when brilliant minds create incredibly intricate ways of

⁴⁶ We could also get to this conclusion by stipulating that Jones does not have a claim to be saved (suppose he was being negligent and got trapped entirely by his own fault). Given that fairness is only concerned with claims, we ought to have a 50-50 lottery between the TV watchers and those watching online. The reason to help Jones would again not be strong enough to outweigh the reason to flip a coin.

consequentialising deontic principles. It strikes me that without side constraints, all consequentialised views, however complex, will inevitably encounter this problem.

Broome also reduces fairness to considerations of proportionality, but fairness in these cases clearly seems to require giving outright priority to the worst off. The death claimant in the Strong Tiebreakers case is not concerned with being given his proportion of the good, but of being given outright priority. Similarly, Jones is not concerned with being given greater weight in our decision, but of being our primary concern. Our judgement that it would be unfair to run a lottery in these cases is based on an understanding of fairness as concerning priority not proportionality. Thus, even if Broome can find a way to capture our judgements, he still seems to do so for the entirely wrong reason.

Lastly, I should note that an important upshot of my argument is that Broome can no longer use his account of fairness to block the inference from cases like Kamm's Sore Throat case to saying that there are irrelevant utilities. If we are going to account for considerations of fairness, we will need some sort of relevance condition. Thus, Broome's choice is to deny that considerations of fairness matter or to admit that there can be irrelevant claims.

Many Tiebreakers (again)

Whilst I think the *Strong Tiebreakers* case is fatal for the outweighing approach, there is one last problem for this approach that I think is worth covering. Broome will happily bite the bullet in this case, so I mention it not to further an argument against Broome, but to help us to better understand how we ought to proceed in these tiebreaking cases. This problem is a version of the Many Tiebreakers case we saw above:

Many Tiebreakers (1): You can save Group A or Group B. In Group A you have a death claim and in Group B you have a death claim and N lost finger claims.

The Outweighing Approach

In the above case, our intuition suggests that the number of lost finger claims matters. If N is small, we hold a lottery. If N is large, we save Group B. This we saw caused a problem for the relevance approach. However, in another version of this case my intuitions differ:

Many Tiebreakers (2): You can save Group A or Group B. In Group A you have a death claim and in Group B you have a death claim and N sore throat claims.

In this version, I share Kamm's intuition (Kamm, 1998, p. 101n3) that no number of sore throats can ever tiebreak between the death claims. Unlike with lost fingers, we should flip a coin no matter the size of N . The sore throat claims are just so weak that taking them into consideration would be disrespectful to the death claimant in Group A.

As before, one might disagree with the specific claim I use here. Perhaps, you think that sore throats can aggregate to tiebreak between competing death claims. Again though, I ask that you consider weaker claims: are there any very weak claims that could not aggregate, no matter how many, to be decisive between death claims?⁴⁷ These could be claims against getting a mild headache or a claim against the brief pain of an injection, or even a claim against a mild itch. All these claims matter, it is not that they never have any weight; they are sufficiently strong that they would give you reason to choose flipping the magic coin over the normal coin when deciding between death claims.⁴⁸ Yet in this case it seems like, no matter how many tiny claims there are, they should not be able to outweigh the reason to run a lottery.

Again, the outweighing approach cannot capture this intuition. According to the outweighing approach, the sore throats are relevant and thus enough of them will outweigh the fairness owed to

⁴⁷ Alternatively, consider the stronger claims being stronger. Could any number of sore throat claims ever tiebreak between which of two planets to save from a death star? Again, the answer seems obviously no, yet this is what is required of any outweighing approach.

⁴⁸ The magic coin will cure the sore throat (or whatever the tiebreaker claim is) whilst still allowing a random choice and thus a fair procedure between the balanced claims (Kamm, 1998, p. 102). Claims that can give reason to choose the magic coin matter at some level and thus do have some weight.

the death claim in Group A, and so we save Group B outright. Broome's approach is thus, once again, overly sensitive to the number of claims in cases where we might think such sensitivity is unintuitive.

Going further

Now I should note that there is no complete consensus even among relevance theorists about the *Many Tiebreakers (2)* case and the intuition in this case is not as strong as it is in the *Strong Tiebreakers* case. As such, it is worth briefly exploring why thought might differ in this case and give a little more argumentation as to why I think we ought to treat the sore throats as unable to tiebreak, no matter the number. This will help to lay the groundwork for the account I develop in the next chapter and explain why my view handles these cases in the way it does.

Firstly, let us look at a version of this case put forward by Victor Tadros, which reports a different conclusion:

Suppose that if I save B there will be a little bit of medicine left which will give me an opportunity to find a cure for the common cold. I will thus save billions of people from suffering sore throats. I then find it intuitive that I ought to save B rather than A. And this seems true even if each person only ever gets one cold, and these colds never result in anyone being harmed beyond the effect of the cold itself. This suggests that the duty-grounding force of C's sore throat is not disabled. (Tadros, 2019, p. 185)

In this example, I must admit that my intuition is with Tadros. However, I am unconvinced that his example is the same as the case above. I think talk of cures obfuscates discussion by bringing in too many other moral considerations. If humanity is cured of sore throats, then, so far as humanity *per se* can have claims, this group claim might be relevant and thus affecting our intuition.⁴⁹ Similarly,

⁴⁹ We might think that by introducing group claims here, I open the door to saying that whilst no single lost finger claim is relevant, the lost finger claimants together have a relevant group claim. Group claims like this might then be enough to explain our intuitions in all the above cases. But this is not the sort of group claim I have in mind when responding to Tadros here. The group claims I have in mind are not combinations of individual claims but of a different sort entirely - they are the claims of a group *as a group*. To distinguish

The Outweighing Approach

our intuition might be affected here by the moral achievement of such a cure. Thus, if we are to address only the issue of aggregating individual claims against harm, we must avoid such aggravating features. In examples where no such features are present, where we can merely treat a very large number (billions even) my intuition remains with running a lottery no matter how many sore throats might be cured.

Once we take out such aggravating features, I think we are in a much better position to properly evaluate our intuitions in these cases.

Furthermore, it is not just my intuition that tells me I should run a lottery at both stages in *Many Tiebreakers (2)*. The reasoning behind the relevance approach in this case is also attractive. As discussed previously, relevance views make appeal to partial concern and the need to consider each person's personal perspective.

In *Many Tiebreakers (2)* then, we must consider the personal perspectives of each of the claimants. Let us first consider the perspective of the dying person in Group A. For her she is about to lose everything. From her perspective the outcome in which she is not saved is "almost like one in which no life would be saved." (Kamm, 1998, p. 167). Naturally the same can be said for the death claimant in Group B, and so the personal perspective of neither is sufficient to determine whom to save.

However, the personal perspective of each death claimant is sufficient to ground a respectful justification for our action to each claimant. We can tell the dying person in Group A that we cannot save her outright as there is another person's life on the line and we must give them a fair chance to be saved too.

between these types of group claims imagine two scenarios. We can save one person's life or 100,000 random people from broken fingers. In this case, we should save the person's life, the broken finger claims are not relevant, individually or as a group. Now suppose, we can save one person's life or the entire population of Tonga (about 100,000 people) from broken fingers. In this case, it seems appropriate to save the broken fingers. Whilst the individual Tongans' claims are irrelevant, the Tongan people as a distinct cultural group have a relevant claim.

Chapter 5 – Two Approaches to Tiebreaking

As discussed in the previous chapter, this respectful justification can be extended further. Respectful justifications can be given to claimants when we cannot give them even a chance of being saved. For instance, we could tell a dying person that we cannot even give her a chance to be saved as we are too busy saving many others from death, or paraplegia etc.

But it seems that sore throats could never form part of this respectful justification to our death claimant in Group A. She only has one life to lead and from her personal perspective she has everything to lose. In that context, we cannot expect her to accept the importance to each individual of having their sore throat cured.

Now consider the personal perspectives of the sore throat claimants. Each of them has a very mild claim: they will be uncomfortable for a couple of days and perhaps they will be unable to enjoy food and struggle to talk for that period. From their perspective, it is a nuisance, an irritant but something that will go away in a week's time, at the latest, and will quickly be forgotten about.

In the context of *Many Tiebreakers (2)*, it would be very easy to give each of them a justification as to why we cannot save them outright. "We are sorry, but we cannot cure your sore throat as there is someone dying and we need to give them an equal opportunity to be saved." A sore throat claimant who pushed back demanding that we take their claim into consideration would quite rightly be called out for acting disrespectfully or unreasonably. To put this in terms of Voorhoeve's permissible personal perspective (Voorhoeve, 2014) a sore throat claimant could not permissibly save herself in this case and leave the death claimant in Group A with no chance at having her life saved. No person with a sore throat can complain if a lottery is run. Not one of them could step forward to ask for their claim to be taken into consideration and so "we retain ties to the personal perspectives of each person" (Kamm, 1998, p. 102).

Furthermore, it strikes me that if we are going to say that claims can be irrelevant when they compete against each other (and, given the Strong Tiebreakers case, we must), then the same

The Outweighing Approach

reasoning that will explain why claims are irrelevant in that realm will also be in play in the tiebreak cases. If the reasoning above works for competing claims, then we need an explanation as to why it would not work for tiebreaking claims.

The only reason I can see that one could forward is that death claims, once matched, lose their power to make other claims irrelevant. But as we saw in the previous chapter, this does not seem right. Matched claimants still have claims of fairness, and these claims of fairness seem to have their own power to make other claims irrelevant. It is not just that it is unfair for the Group A claimant not to be saved, but it is unfair to her to even take into consideration such small claims when her life is on the line. For her to lose an opportunity to be saved over such small claims would be highly disrespectful.

For these reasons it would, at the very least, be more parsimonious if a relevance account played a part of the explanation in sore throat cases – including *Many Tiebreakers (2)*. Nevertheless, as we have already seen, relevance accounts have their own problems. As *Many Tiebreakers (1)* shows, relevance views are entirely insensitive to numbers in tiebreaking cases when we think they ought to be sensitive. We might thus think we have a trade-off; we just have to decide between sensitivity and insensitivity and accept the unintuitive consequences in certain cases.

However, any view that can be variably sensitive and insensitive at the appropriate times will be prima facie stronger. Such a view could capture our intuitions in the *Large Scale, Strong Tiebreaker* and *Many Tiebreaker* cases. I shall develop such a view in the next chapter.

Chapter 6 – Aggregate Relevant Tiebreakers

Introduction

In the previous chapter we saw how both the relevance and outweighing approaches to tiebreaking face difficulties. Both views struggle to capture our intuitions in a variety of different cases. This is because relevance views are insensitive to numbers, whilst outweighing views are overly sensitive to numbers.

In this chapter I develop a hybrid view, *Aggregate Relevant Tiebreakers*, which can get the right answers for the right reasons in all these cases. This view combines the justifications of both the relevance and outweighing approaches to give us a more comprehensive account of fairness when tiebreaking. *Aggregate Relevant Tiebreakers* also avoids the problems with both approaches, draws a compromise between the competing intuitions and can make tiebreaking sensitive to numbers in the right way.

In the first half of the chapter, I will outline what this hybrid account is in simple terms, demonstrate how it captures all the cases, and explain its deeper justification. In the second half, I will combine this hybrid account for tiebreaking with the Sequential Claims Matching by Closeness account developed in Part 1 of this thesis. This should give us a complete account of how we ought to aggregate claims in two option cases. This account will add significant detail and settle a number of ambiguities. However, for reasons of space I cannot address why I decide to develop each element of the view in the specific way that I do. Where necessary, these elements will be noted and addressed in the appendices.

A Simple Solution

In the previous chapter, we saw how both the relevance and outweighing approaches fail in various cases to capture what our intuitions are telling us about these cases, how the relevance approach explains the phenomena better in one case, but worse in others and vice versa. So, it seems we have a dilemma. The way the literature, and by extension I, have treated the two approaches is to treat them as rivals, such that we must decide which approach is preferable.

But there is a very simple solution here. We must recognise that this is a false dilemma. Whilst the two approaches have been presented as rivals, there is no reason that they should not be considered allies.

We can quite easily combine the relevance and outweighing approach into a more comprehensive hybrid account — one that avoids the issues with the relevance approach and the outweighing approach and draws a compromise between the competing intuitions. To do so, I take the relevance approach as the basis of a successful view (it seems like it is closer to the right answer by itself) and amend this view to accommodate an outweighing style approach as per Broome and Hirose. The following is a very simple formulation of the hybrid approach to tiebreaking, that should explain the basic idea:

The hybrid approach: very weak tiebreaking claims can be made irrelevant by the presence of stronger claimants' fairness claims, but relevant tiebreaking claims can still be outweighed by fairness claims.

Fairness claims, in other words, can make other claims irrelevant and have weight of their own. This is already how relevance views treat claims to be saved. Claims to be saved have weight of their own and the ability to make significantly weaker claims irrelevant. That is to say, they have an aggregative and a non-aggregative element. To say that fairness claims have the same properties then is just a very simple extension of this idea.

Now, before we add any detail to this basic idea, it would be helpful to clarify some of the language I will use in this chapter. I distinguish between different types of claims, and the view I develop involves complex interactions and trade-offs between these different claims. The two core types of claims are claims to be saved and fairness claims (i.e., claims to an equal or proportional chance to be saved). To make what follows clearer I shall call these S-claims and F-claims, respectively.

We also have a subset of S-claims which I have been calling tiebreaking claims. It is important to note that tiebreaking claims are not a distinct type of claim with their own particular strength. Instead, a tiebreaking claim is just a role signifier for an S-claim that is not counterbalanced by other S-claims. From now on I shall call tiebreaker claims, T-claims.

With this in view, we can develop a more detailed formulation of the hybrid approach. In homage to Voorhoeve's *Aggregate Relevant Claims* (2014), I shall call this *Aggregate Relevant Tiebreakers*:

Aggregate Relevant Tiebreakers

1. An S-claim is a T-claim if that S-claim is not balanced by a competing S-claim.
2. A claimant has an F-claim if and only if that claimant's S-claim is balanced by a competing S-claim.
3. For the same underlying stake, F-claims are weaker than S-claims. i.e., a claim to be saved from death is stronger than a claim to a fair chance of being saved from death.
4. A T-claim is relevant if and only if it is sufficiently strong relative to the strongest competing F-claim.
5. An F-claim is relevant if and only if it is sufficiently strong relative to the strongest competing T-claim.
6. You should:
 - a. Run a lottery, if and only if, the relevant F-claims outweigh the relevant T-claims.
 - b. Save the group with the greatest sum of strength-weighted, relevant S-claims, if and only if, the relevant T-claims outweigh the relevant F-claims.

A Simple Solution

To see how *Aggregate Relevant Tiebreakers* works, and how it solves the issues outlined, let us go through each of the cases from the previous chapter one by one.⁵⁰ Let us start with Kamm's own case:

Kamm's Large Scale Case: You can save Group A or Group B. In Group A you have 1000 people who will die. In Group B you have 1001 people who will die.

In this case *Aggregate Relevant Tiebreakers* tells us that the 1001st claimant's T-claim is relevant, and so are the F-claims of the competing 1000 claimants: the 1001st person's claim to be saved from death is relevant, as are the other claimant's claims to a fair opportunity to be saved.

Thus, *Aggregate Relevant Tiebreakers* says we must weigh the F-claims against the T-claim. In this case the F-claims outweigh the T-claim and so we should run a lottery. In doing so, *Aggregate Relevant Tiebreakers* captures our intuition, and importantly it does so in the right way.

Aggregate Relevant Tiebreakers also handles well *Strong Tiebreakers*:

Strong Tiebreakers Case: You can save Group A or Group B. In Group A you have N number of sore throat claims. In Group B you have N number of sore throat claims and a death claim.

In this case *Aggregate Relevant Tiebreakers* says that the sore throat claimant's F-claims are not relevant to the death claimant's T-claim. Thus, no number of sore throats can ever outweigh the death claim - and so we save Group B outright. Once again, this gets the right answer.

Lastly let us look at the two *Many Tiebreakers* cases together:

⁵⁰ It should be noted, that for simplicity, I have developed *Aggregate Relevant Tiebreakers* here (and the Complete Account below) such that we weigh F-claims against T-claims directly. Nevertheless, a more complex statement of *Aggregate Relevant Tiebreakers* could adjust the weight of the F-claims and T-claims according to expected value. This would factor in the chance that running a lottery will involve saving the larger group anyway. See Appendix 10, for the more complex statement and why adjusting for expected value should make no difference.

Many Tiebreakers (1): You can save Group A or Group B. In Group A you have a death claim and in Group B you have a death claim and N lost finger claims.

Many Tiebreakers (2): You can save Group A or Group B. In Group A you have a death claim and in Group B you have a death claim and N sore throat claims.

In the first of these cases, *Aggregate Relevant Tiebreakers* tells us to run a lottery if N is small and save Group B if N is large. This is because the lost finger claimants' T-claims are relevant to the death claimant's F-claim. However, if N is small the death claimant's F-claim will outweigh the lost finger claimants' T-claims. Conversely, if N is large, the T-claims will outweigh the F-claim. Thus, *Aggregate Relevant Tiebreakers* is appropriately sensitive to the numbers in this case.

In the second case, *Aggregate Relevant Tiebreakers* tells us to always run a lottery. This is because the sore throat claimants' T-claims are never relevant to the death claimant's F-claims. Thus, the sore throats are dismissed from consideration before they can be weighed against the considerations of fairness. Again then, *Aggregate Relevant Tiebreakers* captures our intuitions and explains why our judgements are not sensitive to numbers in this case.

Importantly, the way *Aggregate Relevant Tiebreakers* handles the *Many Tiebreakers* cases explains why our intuitions seem to conflict in the two cases. It allows us to capture both intuitions, which we might otherwise think are inconsistent. Where we draw the line for relevance is obviously up for debate, but at least by untangling these two intuitions we might more accurately be able to decide where this line might be.⁵¹

⁵¹ A careful reader might have noticed a different potential problem for this account; *Aggregate Relevant Tiebreakers* seems to violate the *Principle of Net Addition* discussed in Chapter 1. Namely, that adding a stronger group of claims to an already winning group of claims, will make that winning group of claims less choice worthy. If so, this would pose a significant problem to my account. For reasons of space, I do not address this objection here, for more detail see Appendix 9.

Deeper Justification

So, *Aggregate Relevant Tiebreakers* can capture what we want to say in all these cases, but that is only half of the story. We need to understand the motivation of the view and its underlying rationale if *Aggregate Relevant Tiebreakers* is to be fully acceptable. Firstly, I think it is important to note that the fact that *Aggregate Relevant Tiebreakers* can capture our intuitions in all the above cases is at the very least prima facie reason to trust *Aggregate Relevant Tiebreakers*. I follow Kamm (Kamm, 1998, pp. 5-9) in putting quite a lot of trust in intuitions - we can trust our intuition or pre-theoretical judgements unless we have good reason to think otherwise.

A reason we might have not to trust our intuitions in a particular area is that the intuitions clash or are contradictory. Motivating a view or finding some deeper rationale is much more important in cases where our intuitions clash and where we might have to abandon one or more of our pre-theoretic judgements - this is because we need to find a reason to prefer one intuition over the other. However, in cases where we can capture all our intuitions or our intuitions do not clash then, whilst it remains important to understand the underlying motivation or deeper rationale, the rationale is there to explain the intuitions rather than judge between the intuitions or justify the view.

Thus, where *Aggregate Relevant Tiebreakers* can capture all our intuitions, I believe there is less need for a deeper rationale to justify the view. The fact that *Aggregate Relevant Tiebreakers* can capture all our intuitions is itself good reason to believe that *Aggregate Relevant Tiebreakers* is a good explanation of the phenomena. *Aggregate Relevant Tiebreakers* vindicates our intuitions, and so I believe the explanatory burden for why we should not trust our intuitions shifts to those views that have to bite bullets or explain away some of our pre-theoretic judgements.

Now, I have shown, through *Aggregate Relevant Tiebreakers*, how the relevance and outweighing approaches are compatible. However, I have not shown how the deeper rationales of these two approaches are compatible. One might think that my view, as a hybrid view, inherits the explanatory

burden of both views individually, and the burden of showing how the two deeper rationales fit together. Yet, part of the explanatory burden of a relevance only approach is that we need to explain why, contra to common sense morality, proportionality does not matter. Similarly, Broome's outweighing approach needs to explain why priority does not matter. Plus, both views need to explain why our intuitions are wrong in some of these cases. The hybrid view has none of these problems. The theory might be more complex, but it is explanatorily simpler as it captures all the phenomena and avoids the problems of the views taken separately. As Ross puts it: "[I]oyalty to the facts is worth more than a symmetrical architectonic or a hastily reached simplicity." (Ross, 1930, p. 23)

Nevertheless, let us not shirk our responsibility here. I recognise that not everybody puts as much stock in their intuitions as Kamm and I. Thus, it would be helpful to give a sketch of the underlying rationale of the view. After all, our intuitions are picking up on some underlying structure and so explaining what our intuitions tell us about our ethical principles or reasons is still going to be helpful.

Now, there are detailed and impressive arguments for both approaches which I will not review here.⁵² Instead, I shall focus on how the two deeper rationales are compatible. My aim is not to defend both approaches, but just to defend their compatibility. Nevertheless, to do so we do need to understand the underlying rationales of the two views. As we saw in the previous chapter, the relevance approach can be characterised by the following claims about fairness:

1. Fairness is about side constraints; it rules out certain considerations. (Kamm, 2002, p. 686)
(Voorhoeve, 2014, p. 70)
2. Fairness involves giving "*significant priority* to those who are the worst off" (Kamm, 2002, pp. 686-687). (Voorhoeve, 2014, pp. 72-75)

⁵² In particular, see Kamm (1998, 2002, 2007) and Broome (1990, 1998, 2002, 2004) respectively.

Conversely, we saw that the outweighing approach is characterised by the following claims:

1. Fairness is an end state, and only provides pro tanto reasons. (Broome, 1998, p. 2) (Hirose, 2015, pp. 192-195)
2. Fairness requires “that *claims should be satisfied in proportion to their strength.*” (Broome, 1990, p. 95) (Hirose, 2015, pp. 189-190)

Once we see this, it seems that the two approaches are at complete odds. If so, this spells disaster for the hybrid view. However, as we have seen in the previous chapter, both of these approaches have serious shortcomings. Our intuitions in several cases seem to suggest that sometimes fairness works as a side constraint and sometimes as a valuable end state. This is because side constraints are not sensitive to numbers whilst end states are sensitive to numbers and, intuitively, we want fairness concerns to sometimes be sensitive to numbers and sometimes not.

For instance, in *Many Tiebreakers (1)*, we want fairness to work as an end state, such that lost fingers can outweigh considerations of fairness if they are strong enough. Conversely, in *Many Tiebreakers (2)*, we want fairness to work as a side constraint and prevent any number of sore throats from deciding between the two death claims. So ideally, to capture all the moral phenomena, we want a view of fairness that provides both side constraints and end state value.

But the literature has so far treated this as a choice we have to make - we can either treat fairness as a side constraint or an end state, but not both. I think there are a couple of reasons why fairness has been treated this way. Firstly, there is a question as to how fairness could play two such different roles. How could it constrain the reasons provided by other features and provide reasons of its own? We might even think doing so would double count considerations of fairness. Secondly, how could we distinguish between when fairness acts as a side constraint and when it acts an end state? If we choose to treat fairness as a side constraint or as an end state when it suits us, then this seems like an ad hoc and arbitrary solution.

This is where it becomes important to recognise the other ways the two approaches characterise fairness. Kamm and other relevance theorists characterise fairness as being concerned with priority and Broome characterises fairness as being concerned with proportionality. Again, these are treated as rival views.

But I just do not see why these have to be understood as competing understandings of what fairness requires. Intuition and common-sense morality suggest that proportionality and priority are both aspects of fairness that need to be accommodated for if we are to have a full understanding of fairness. A simple case illustrates this point.

Suppose I have two children who are playing football. In one particularly bad tackle they both end up hurt. My daughter scrapes her knee, but my son has obviously broken his ankle. Now suppose I attend to my daughter first; I clean up her wound and put a plaster on her knee. Only then do I attend to my son; I take him to the hospital and spend hours with him in A&E until they put a cast on him. Both my children have got what they proportionally deserve, but my son would still have a good claim to say that I treated him unfairly. As he was hurt much worse, I should have attended to him first - I should have given him priority.

Conversely, suppose that I do attend to my son first. But when we return from the hospital, I do not even check up on my daughter to see how her knee is. In this case, not only am I a bad parent, but my daughter has a claim that I have treated her unfairly. By attending to my son properly but not her, I have not satisfied her claim to the proportion of its strength.

Thus, it is fairly easy to see that fairness involves both considerations of priority and of proportionality. Now, I do not want to say that this is all fairness involves, there may be further features of fairness, but here I am only interested in priority and proportionality. Fairness, then, might be best understood as an umbrella concept for a number of different, but related, moral considerations.

Deeper Justification

Once we recognise this, we can also see how fairness can be both a side constraint and an end state. There are two underlying elements of fairness, which can play the two different roles of fairness. Considerations of priority work as side constraints, whilst considerations of proportionality are end states. This gives us the ability to non-arbitrarily determine when fairness works as a side constraint and when it works as an end state. It also explains why fairness is not being double counted - we are counting two different features, sometimes both features will be present and sometimes only one feature will be present.

Thus, priority explains the anti-aggregative element of fairness; if a claimant has priority over others, then those other claims cannot be aggregated against the priority claim. Whilst proportionality explains the aggregative element of fairness; the claims to proportional treatment can add up and outweigh (or be outweighed by) other relevant claims. For instance, in the Scanlon-Style World Cup case, Jones has priority over the world cup watchers, and so we do not aggregate the world cup watchers' claims against Jones' claim. Whilst in the Large-Scale case the death claimants do not have priority over each other, but they do have claims to be treated equally. The claims of each of the 1000 dying people to be treated equally to the 1001 can thus aggregate and outweigh the extra life in the larger group.

A hybrid view becomes quite intuitive once we think about it in this way. It captures more of what we want to say about fairness than either view by itself does, and it better fits with the limited aggregation thesis that sometimes aggregation is permitted and sometimes it is not. Plus, as *Aggregate Relevant Tiebreakers* shows, it captures more of the moral phenomena. Thus, whilst we might have thought that the hybrid approach would involve taking on the explanatory burden of both views, and so would be explanatorily more complex, the reality is that each approach covers the gaps of the other approach.

Part of the explanatory burden of the relevance approach is that we need to explain why, contra to common sense morality, fairness is not valuable in and of itself. Conversely, the outweighing

approach needs to explain why fairness is only an end state value and why priority does not matter. Plus, both views need to explain why our intuitions are wrong in some of these cases. The hybrid view has none of these issues. The theory might be more complex, but it is explanatorily simpler as it captures all the phenomena and avoids the problems of the views taken separately.

Furthermore, the hybrid approach also explains why our intuitions about fairness seemed to conflict. Concepts of priority and proportionality both have merit and so our intuitions are not confused but are picking up on different morally important aspects. By recognising that we can maintain both at the same time and by distinguishing between the two types of fairness we untangle these intuitions. In doing so we dissolve the conflict, and we make taking fairness considerations into account much easier.

The fact that considerations of proportionality and priority are so closely linked in these cases, and work in tandem, might also go some way to explain why we have become so confused about them and have believed them to be rival understandings of fairness. They are so close to each other, and so interconnected, that it is not surprising that we have treated them together as fairness for so long. Hopefully, in prising the two concepts apart, and then putting them back together again, we will have a better understanding of fairness as a whole, and our duties in these particular cases.

A Complete Account of Aggregation

I now want to synthesise the different elements of my aggregation view into one cohesive account. This means combining *Aggregate Relevant Tiebreakers* with *Sequential Claims Matching by Closeness*, in a way that captures everything we want to say about aggregation in complex cases.

To do so I first need to amend *Sequential Claims Matching by Closeness* such that *Aggregate Relevant Tiebreakers* is compatible with it. Thankfully, this amendment also helps to simplify *Sequential Claims Matching by Closeness*. This will lay the groundwork for the rest of the account.

A Complete Account of Aggregation

I then need to add some detail to *Aggregate Relevant Tiebreakers*. *Aggregate Relevant Tiebreakers* as I have developed it above works well in homogenous and simple cases. However, it is unclear how *Aggregate Relevant Tiebreakers* works in more complex and heterogenous cases with lots of different F-claims and T-claims. I shall add a number of steps to *Sequential Claims Matching by Closeness* that accurately capture *Aggregate Relevant Tiebreakers* and give guidance on how we weigh F-claims and T-claims in these more complex cases.

For reasons of space, I shall not explain in great detail why I develop these steps in the way I do. There are a number of different ways we could develop *Aggregate Relevant Tiebreakers* in these heterogenous cases, and how we combine it with *Sequential Claims Matching by Closeness*. I shall simply state how I have developed the view, and signpost towards a number of appendices that explain my reasoning in these cases and the rival options.

Lastly, I will use a complex test case to demonstrate how the new sequence works, and how it can calculate our all considered reasons in highly complex cases.

A Note on Notation

However, before we continue, I need to introduce some new notation that shall help to make the following, quite complicated, cases easier to follow. It is easiest to see this notation in practice so let us consider *Kamm's Large Scale Case*:

Kamm's Large Scale Case

Level	Group A	Group B
1 - Death	1000	1001

In this case, we have one thousand dying people in Group A and one thousand and one dying people in Group B. The first thing we do, following *Sequential Claims Matching by Closeness*, is to match an equal number of S-claims from each side until one group has no more S-claims. To indicate this, we strikethrough the matched claims, and put any remaining S-claims in square brackets:

Kamm's Large Scale Case

Level	Group A	Group B
1 - Death	1000	1001 [1]

Now, once we match the death claimants together, they no longer have S-claims, but they do have F-claims. To capture these F-claims in the table we shall put the F-claims in separate angle brackets <>:

Kamm's Large Scale Case

Level	Group A	Group B
1 - Death	1000 <1000>	1001 [1] <1000>

Now we can more easily see that we have 1000 F-claims in each group and one T-claim in Group B. Lastly, then to decide whether we ought to run a lottery or save Group B outright we need to weigh the T-claim against the competing F-claims. For illustrative purposes, let us assume that ten F-claims match with exactly one T-claim on the same level. To indicate this on our table we can again strikethrough the claims in each group:

Kamm's Large Scale Case

Level	Group A	Group B
1 - Death	1000 <<1000> <990>	1001 [1] <1000>

After striking through these claims, we can see that there are no T-claims left; only F-claims remain. Thus, we must run a lottery. Now, we might think there is an imbalance here, as Group B has more unmatched F-claims than Group A. However, this does not matter as F-claims are not in competition with each other because they are claims to the same outcome - i.e., they are all claims to a fair opportunity of being saved.

Hopefully, this notation should be clear, and will help with some of the more complex cases in this section. Whilst I have walked through each step very slowly in this example, do bear in mind that for the sake of brevity I shall skip most of these steps in the following cases.

Amending Sequential Claims Matching by Closeness

Firstly, before we amend and precisify *Aggregate Relevant Tiebreakers*, we first need to amend *Sequential Claims Matching by Closeness*. It is important to note that, F-claims (and thus *Aggregate Relevant Tiebreakers*) only come into play when all S-claims that can be matched have been matched. In short, this is because we need to work out how the two groups compare before we consider tiebreaking between them.⁵³ So we must match all the S-claims first before we consider T-claims and F-claims. In Chapter 3 I develop this sequence to capture how we should match S-claims:

Sequential Claims Matching by Closeness

- I. Identify the strongest claim still in consideration and ask, does it have *anchoring* relevance?
 - a. If so, go to Step II.
 - b. If not, then it is not the case that you should decide in favour of one group over the other (though you must save one).
- II. Does the other group contain claims of types that are of *competing* relevance to this claim and sufficient in number to match it?
 - a. If so, go to Step III.
 - b. If not, decide in favour of the group with the strongest claim still in consideration.
- III. Identify the strongest remaining claim *with* competing claims [zero] levels lower. Match these claims and eliminate them from consideration.
- IV. Continue step III until there are no claims with competing claims [zero] levels lower.
- V. Repeat steps I - IV, with competing claims one level lower, then two levels lower, and so on.
- VI. Continue until either:
 - a. One group contains unmatched claims with anchoring relevance, 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).

⁵³ I address the reasoning here in more detail in Appendix 3.

However, this sequence will not work with *Aggregate Relevant Tiebreakers*. Step II and VI will tell us to tiebreak without comparing the strength of the T-claims to the relevant F-claims. This view currently captures a relevance only approach to tiebreakers.

Thus, I suggest we simplify *Sequential Claims Matching by Closeness* significantly to remove these aspects of the sequence, such that we can address them with *Aggregate Relevant Tiebreakers* instead. Furthermore, I will remove Step I to further simplify the sequence as this step can also be captured as part of *Aggregate Relevant Tiebreakers*.

Simplified Sequential Claims Matching by Closeness

- I. Identify the strongest remaining S-claim *with* competing S-claims [zero] levels lower. Match these S-claims and eliminate them from consideration
- II. Repeat step I until there are no S-claims with competing S-claims [zero] levels lower.
- III. Repeat steps I - II, with competing S-claims one level lower, then two levels lower, and so on until no S-claims with relevant competing S-claims remain.

This leaves us with a much more basic version of *Sequential Claims Matching by Closeness*, which captures how we should match S-claims with each other. It follows the *Match by Closeness* principle but does not tell us how to decide between the matched claims. Importantly, it also means that we are matching all the S-claims before we match any F-claims. We can then build on this by adding *Aggregate Relevant Tiebreakers* to the end of this sequence to decide how to handle the T-claims and F-claims.

Amending Aggregate Relevant Tiebreakers

With this in view, we can move on to finally amending *Aggregate Relevant Tiebreakers* and introducing it to our sequence. To do so we shall introduce one step of *Aggregate Relevant Tiebreakers* at a time. I shall briefly note why I introduce each step as I do as each step is developed with particular principles and cases in mind, and in response to possible ambiguities. There is simply

not space for me to go into the detail of my reasoning for each of the steps here. But I address this reasoning in a series of appendices, which I will signpost at each stage.

The first step that we will want to introduce to our sequence is the following step:

- IV. Consider the remaining S-claims. Dismiss all S-claims that are not relevant to the strongest F-claims. I.e., dismiss all S-claims that do not have anchoring relevance.

This step captures the role of anchoring relevance, as discussed in Chapter 4: irrelevant T-claims can be dismissed from consideration early because T-claims need anchoring relevance to be able to tiebreak between balanced S-claims. In other words, we dismiss all potential T-claims from consideration if they could not form part of a justification not to give the claimants with the strongest S-claims a fair chance to be saved. Thus, it is this step that captures the relevance aspect of my tiebreaking view.

The second step we need to introduce is the following:

- V. Dismiss from consideration any T-claims that are competing irrelevant to any other competing T-claims.

This step is needed to prevent an issue with competing T-claims⁵⁴. There is a possibility that we could have competing T-claims in both groups if the weaker T-claims are competing irrelevant to the stronger T-claims. This happens because the weaker T-claims are not taken out of consideration by the matching process. Thus, *Aggregate Relevant Tiebreakers* needs a step that eliminates the weaker T-claims that are irrelevant to stronger T-claims, such that they cannot affect our decision. Step V fulfils this role and allows us to maintain that significantly weaker S-claims cannot form part of a justification not to save significantly stronger S-claims.

⁵⁴ For more detail on this, see Appendix 5.

With these easier steps out of the way we can move on to the more complicated issue of how to capture the way we match F-claims with T-claims. Unlike with T-claims, we cannot simply dismiss all the irrelevant F-claims in a couple of steps. F-claims only need to be locally relevant to the competing T-claims. I cover why in more detail in Appendix 6. In short however, this is because F-claims do not become the anchoring claims. Their only role is to outweigh the T-claims and therefore they only need competing relevance and, as we saw in Chapter 4, competing relevance is determined by local relevance. As such F-claims might be relevant or irrelevant depending on which T-claims have already been matched and taken out of consideration. Thus, we will need a more complex sequence, like that of *Sequential Claims Matching by Closeness*, to capture this aspect of *Aggregate Relevant Tiebreakers*.

Now, we will want to match the T-claims and F-claims following either *Match by Closeness* or *Match in Order of Strength*. I prefer *Match by Closeness* because it better fits with the reasoning provided in the rest of the thesis and so I will amend *Aggregate Relevant Tiebreakers* in this way. However, there might be substantial disagreement on this point, and it is not entirely clear that one approach is definitively better than the other. I address this discussion in significantly more detail in Appendix 4 where I have also created a version of this sequence consistent with *Match in Order of Strength* when comparing T-claims and F-claims. Whilst we might disagree about the process here, it should be noted that both matching principles give the same answer eventually.

Now, to capture the *Match by Closeness* approach to *Aggregate Relevant Tiebreakers* let us introduce the steps slowly. The first step is the following:

- VI. Identify the strongest remaining T-claim *with* relevant competing F-claims [zero] levels higher (i.e., claims with the same amount at stake). Match these claims and eliminate them from consideration.

Steps VI is the *Aggregate Relevant Tiebreakers* equivalent of Step I of *Sequential Claims Matching by Closeness*, by which I mean they play the same role in the sequence. Step VI identifies claimants with

A Complete Account of Aggregation

the same amount at stake but (as opposed to Step I) with different types of claims (F-claims and T-claims) and will compare and match their respective claims. To make this clear consider the following case:

Case 15

Level	Group A	Group B
1 - Death	100 <100>	101 [1] <100>
2 - Paraplegia	100 <100>	101 [1] <100>
3 - Lost Limb		1
4 - Lost Hand		

Step VI will identify the death claimant with a T-claim in Group B as having a claim on the same level as the death claimants with F-claims in Group A and thus will match these claims first. The death claimant's T-claim is more relevant to the other death claimants' F-claims than the paraplegia claimant's T-claim is. Now, this is true despite the fact that a paraplegia claimant's T-claim is stronger than a death claimant's F-claim, thus relevance and strength come apart.⁵⁵ Of course, the strength of the death claimant's T-claim is stronger than the other death claimant's F-claims, and thus it will take several F-claims to match the T-claim.

Step VI is followed by Step VII, which, much like Step II, simply tells us to repeat Step VI:

VII. Repeat Step VI until there are no T-claims with competing F-claims [zero] levels higher.

In Case 15, after the death claimant's T-claim is matched, we compare the paraplegia claimant's T-claim to the paraplegia claimants' F-claims in Group A. Again, we compare the paraplegia claimant's T-claim to the paraplegia claimants' F-claims because they have the same underlying stake. We could compare this T-claim to any remaining death claimants' F-claims – and doing so would get the same answer (see Appendix 4) – but it seems more appropriate to compare it to the F-claims of those with the same amount at stake.

⁵⁵ See Appendix 4, especially the latter part on *Weight Adjusted Match by Closeness* for more on this.

Together then, Steps VI and VII consider all T-claims with competing F-claims of those on the same level (i.e., with the same amount at stake) and match those claims. Thus, they match all the closest claims first.

Step VIII completes the matching part of the sequence. Again, this is very similar to Step III of *Sequential Claims Matching by Closeness* above. Step VIII, simply tells us to repeat Steps VI and VII with claims one level apart, then two levels apart:

VIII. Repeat steps VI- VII, with competing F-claims one level higher, then two levels higher, and so on.

In Case 15, this involves identifying the lost limb claimant's T-claim as being only one level below the paraplegia claimants' F-claims and so matching these claims together. This means the sequence carries on matching by closeness, as we compare slightly less similar T-claims and F-claims with each cycle through the sequence.

One element of this Step that you might have noticed is different is that Step VIII, unlike Step III, tells us to compare T-claims with the closest competing F-claims on increasingly higher levels, as opposed to lower levels. Now, I address why I develop the view in this way in much more detail in Appendix 7, but the short version is that F-claims of weaker underlying S-claims should always be considered irrelevant to stronger T-claims. I call this principle *Stronger Always Tiebreaks*:

Stronger Always Tiebreaks: A T-claim always tiebreaks between F-claims of those with less at stake. I.e., F-claims are irrelevant to the T-claims of claimants with more at stake.

In other words, a stronger claim will always tiebreaker between equally balanced weaker claims. For instance, if we can save equal numbers of people from paraplegia, but if we save the second group, we can also save a life, we should always save the second group. I come to this conclusion as the result of a process of reflective equilibrium. It seems intuitive that in cases like the above, the death claim should always tiebreak between the paraplegia claims. There also seem to be good theoretical

reasons to prefer the approach – namely that it seems to better respect the personal perspectives of the claimants and gives appropriate priority to the stronger claimants.

Furthermore, even if we reject this approach, we will still need to provide a limit of relevance at which F-claims cannot be aggregated against stronger claimants' S-claims. Otherwise, we will end up with cases, such as *Strong Tiebreakers* from the previous chapter, where sore throat claimants' F-claims can be aggregated against a death claim. Saying that F-claims stop being relevant when the competing T-claims are stronger seems like a natural place to put the limit, and at the very least as good as any other option.

Lastly Step IX sets out the 'win' conditions when all the T-claims and F-claims are finally matched:

- IX. Continue until either:
- a) there is a T-claim that cannot be outweighed by F-claims, in which case you should decide in favour of that group; or
 - b) there are no T-claims left, in which case use a fair procedure.

This step should be fairly self-explanatory. If, after matching F-claims and T-claims, there are unmatched T-claims then we save that group. If all T-claims are taken out of consideration, then there are no claims which can tiebreak, so we run a lottery. This step settles all possible cases without any ambiguity.

Now, this account does mean that if T-claims and F-claims are matched exactly we proceed by running a fair lottery. It thus treats the fair procedure as default and that we need an active reason to choose one group outright. This strikes me as the right approach, given that a fair procedure gives everyone a chance whilst saving one group outright condemns some claimants to not having a chance of being saved.

I recognise that I have not provided much of an argument here, but I do not think there is much of an upshot of doing so. None of the rest of my view rides on how we handle such cases. Thus, if one

has a strong intuition, argument or principle that explains how we should settle cases when F-claims and T-claims are precisely balanced, they are welcome to substitute that approach into the account. Step IX can be easily amended to accommodate several different approaches to such cases.

The Complete Account

To put it all together we have the following complex, but clear sequence:

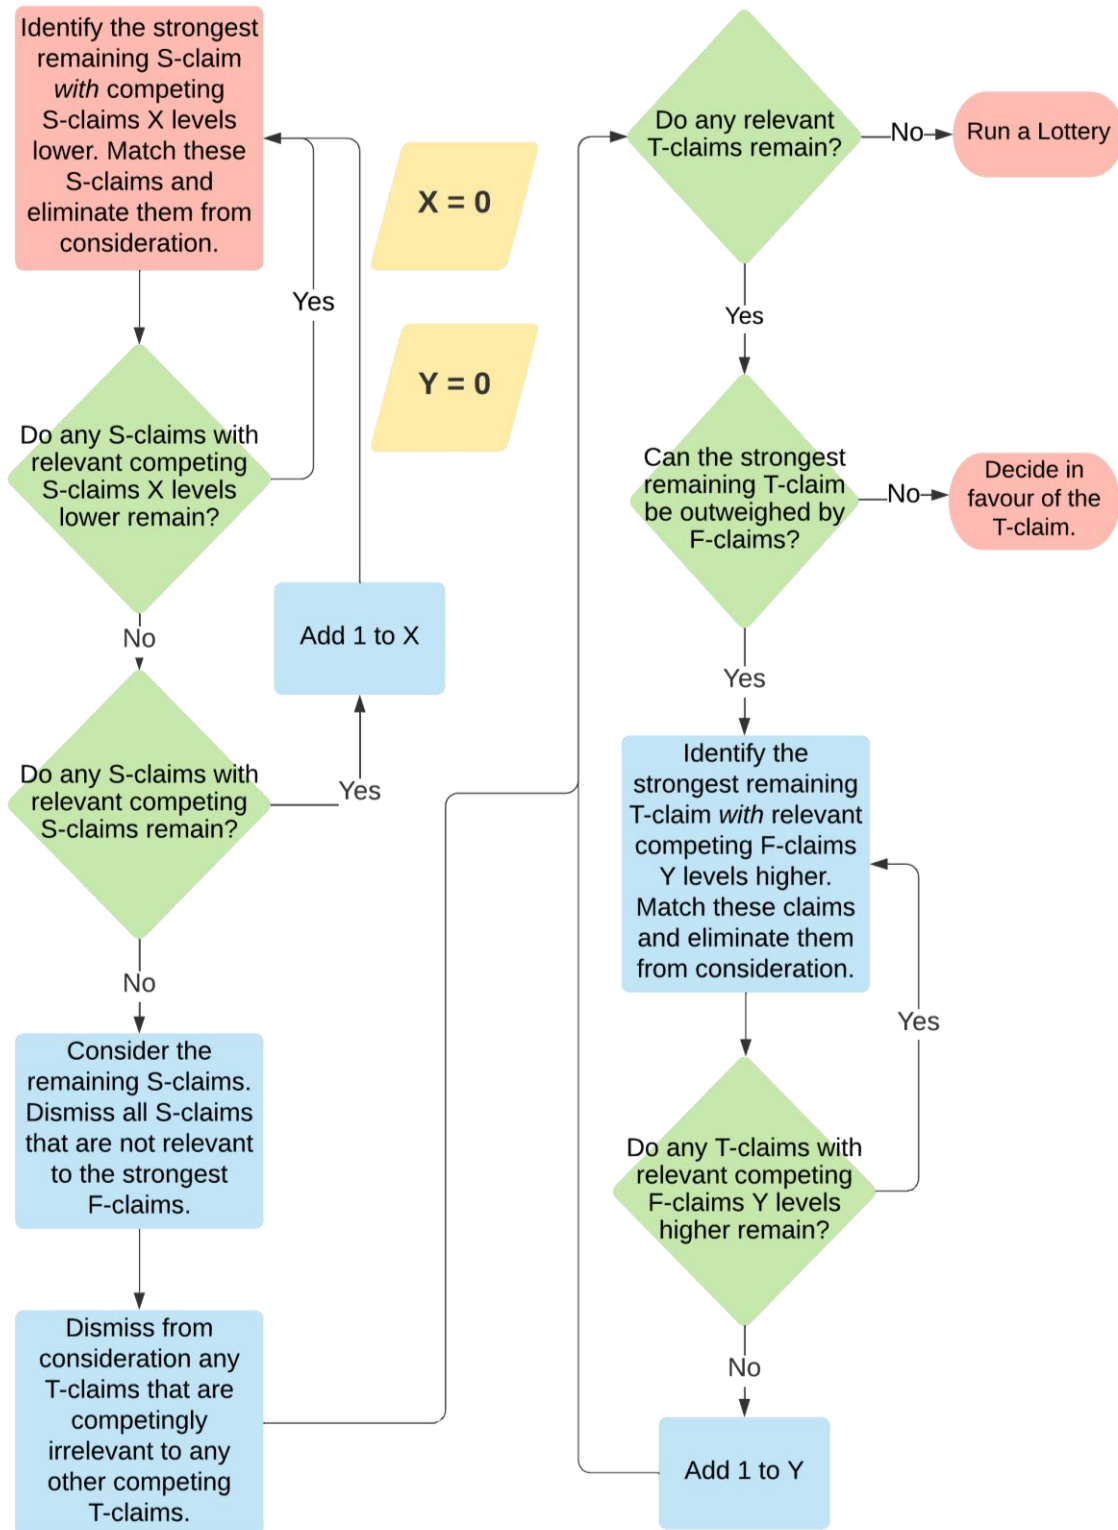
The Complete Account

- I. Identify the strongest remaining S-claim *with* competing S-claims [zero] levels lower. Match these S-claims and eliminate them from consideration.
- II. Repeat Step I until there are no S-claims with competing S-claims [zero] levels lower.
- III. Repeat Steps I - II, with competing S-claims one level lower, then two levels lower, and so on until no S-claims with relevant competing S-claims remain.
- IV. Consider the remaining S-claims. Dismiss all S-claims that are not relevant to the strongest F-claims. I.e., dismiss all S-claims that do not have anchoring relevance.
- V. Dismiss from consideration any T-claims that are competingly irrelevant to any other competing T-claims.
- VI. Identify the strongest remaining T-claim *with* relevant competing F-claims [zero] levels higher (i.e., claims with the same amount at stake). Match these claims and eliminate them from consideration.
- VII. Repeat Step VI until there are no T-claims with competing F-claims [zero] levels higher.
- VIII. Repeat Steps VI- VII, with competing F-claims one level higher, then two levels higher, etc.
- IX. Continue until either:
 - a) there is a T-claim that cannot be outweighed by F-claims, in which case you should decide in favour of that group; or
 - b) there are no T-claims left, in which case use a fair procedure.

A Complete Account of Aggregation

This sequence captures how we ought to match S-claims with each other, and how we should match left over S-claims (i.e., T-claims) with the F-claims of those whose S-claims have been considered and balanced. As before, I include a flow chart for increased clarity:

The Complete Account - Flow Chart



A Test Case

To make the whole sequence easier to understand, let us walk through a test case one step at a time. Again, let us assume that competing relevance extends two levels and anchoring relevance extends four levels. In other words, an S-claim is relevant to an S-claim two levels higher, whilst a T-claim is relevant to an F-claim 4 levels higher.

To see how this sequence can help us in particularly complex cases where it is really hard to rely on individual moral judgement or our intuitions in that particular case, I have made our test case particularly unwieldy:

A Complex Test Case

Level	Group A	Group B
1 - Death	8	11
2 - Paraplegia		
3 - Lost Limb	10	
4 - Lost Hand		20
5 - Lost Finger	10	
6 - Broken Arm	100	

Step I (Round 1): For the first step we identify the strongest S-claims with competing S-claims on the same level; these, of course, are the death claims so we match these claims and take them out of consideration.

A Complex Test Case

Level	Group A	Group B
1 - Death	8 <8>	11 [3] <8>
2 - Paraplegia		
3 - Lost Limb	10	
4 - Lost Hand		20
5 - Lost Finger	10	
6 - Broken Arm	100	

Step II (Round 1): We repeat Step I with any other S-claims on the same level, but there are no more S-claims on the same level, so we skip to Step III

Step III: We repeat Steps I and II but with S-claims one level apart:

A Complete Account of Aggregation

Step I (Round 2): We identify the strongest S-claims with competing S-claims one level lower. These are the lost limb and lost hand claims. We match these S-claims and take them out of consideration. (We have stipulated that two S-claims match with exactly one S-claim on the tier above)

Step II (Round 2): Again, there are no more S-claims with competing S-claims one level lower, so we skip to Step III.

Step III (Round 2): Would tell us to repeat Steps I and II again with competing S-claims two levels lower. However, there are no remaining S-claims with competing claims two levels lower.

After completing Step III and the *Sequential Claims Matching by Closeness* (i.e., the non-tiebreaking) aspect of the sequence our case looks like this:

A Complex Test Case

Level	Group A	Group B
1 - Death	8 <8>	11 [3] <8>
2 - Paraplegia		
3 - Lost Limb	10 <10>	
4 - Lost Hand		20 <20>
5 - Lost Finger	10	
6 - Broken Arm	100	

Step IV: We next eliminate from consideration any S-claims that do not have anchoring relevance. In other words we dismiss any S-claims that is not relevant to the strongest F-claim with which they compete. We stipulated that anchoring relevance extends four levels, so the Broken Arm S-claims do not have anchoring relevance and so are dismissed from consideration.

Step V: We next dismiss from consideration any T-claims that are not competently relevant to the strongest T-claims. Here we can see that the lost finger T-claims, whilst anchoringly relevant, are not competently relevant to the three death T-claims in Group B. Thus, we also dismiss from consideration the lost finger T-claims.

A Complex Test Case

Level	Group A	Group B
1 - Death	8 <8>	11 [3] <8>
2 - Paraplegia		
3 - Lost Limb	10 <10>	
4 - Lost Hand		20 <20>
5 - Lost Finger	10	
6 - Broken Arm	100	

Step VI: We next identify the strongest relevant T-claim with competing F-claims and match those claims. These of course are the three death T-claims in Group B and the eight death F-claims in Group A. Let us stipulate that four F-claims exactly match with a single T-claim of the same level such that all eight F-claims counterbalance two of the T-claims. This gives us the following:

A Complex Test Case

Level	Group A	Group B
1 - Death	8 <8>	11 [3] [1] <8>
2 - Paraplegia		
3 - Lost Limb	10 <10>	
4 - Lost Hand		20 <20>
5 - Lost Finger	10	
6 - Broken Arm	100	

Steps VII and VIII: Both these steps tell us to match other T-claims with F-claims on the same level and then with F-claims on progressively higher levels. However, there are no more relevant competing F-claims. This is because, as we saw above and as I address in more detail in Appendix 7, F-claims with weaker underlying S-claims are never relevant to stronger T-claims. In this case, the lost limb claimant’s F-claims are not relevant to the dying person’s T-claim. Thus, we have completed all possible matchings, and the remaining F-claims are dismissed from our consideration.

Step IX: The last step recognises that the last remaining T-claim cannot be outweighed by competing F-claims, and so concludes that we ought to save Group B.

Hopefully, seeing the sequence in action should make clearer exactly how the sequence works.

Though it looks complex, in reality, when it is broken down and once practiced, it is fairly easy to

implement. The sequence breaks down what is a very complicated case into much more manageable parts and allows us to come to a well-reasoned conclusion. Rather than have an overall intuition about what we ought to do in a case like this (which I believe is far too complicated to have accurate intuitions about anyway), we can break the case up into smaller competitions which we can have accurate intuitions about and then have a way to bring all these features together to give an overall decision on what we ought to do all things considered. It also much more accurately captures the complex way in which the various features in these cases interact. It captures the holistic nature of these cases in a calculable way.

To conclude this chapter, we have seen how my hybrid account can capture all our intuitions in these tiebreaking cases. We have also seen that this is not an ad hoc solution; it is not simply a way to capture our intuitions without deeper normative justification. The reasoning I provide for my Hybrid view matches the reasoning already present in many relevance views when it comes to S-claims. *Aggregate Relevant Tiebreakers* just expands this reasoning to F-claims and T-claims. Thus, saying that T-claims must be relevant *and* sufficiently strong on aggregate is not adding any complexity that does not already exist on most limited aggregation views. Saying that F-claims can both make other S-claims irrelevant and provide weight of their own is just to say that F-claims act in the same way as S-claims!

Furthermore, once we recognise that the relevance and outweighing approaches are not rival approaches, we can also see that the underlying normative reasons behind both accounts are not at odds either. Thus, *Aggregate Relevant Tiebreakers* does not need some brand-new basis on which to stand but can rely on the reasoning provided for the relevance and outweighing approaches, and thus provides a more comprehensive account of what we ought to do, all things considered.

Conclusion

This thesis is long enough already so I shall not draw it out with a long summary and conclusion.

However, I do want to draw attention to some issues that will need covering elsewhere and a few of the upshots of the view to bring discussion full circle.

Let us start with a couple of remaining issues. The first is simply to state that in the thesis I only cover two option cases, where we can choose between Group A and Group B. Of course, in reality we are very rarely left with only two options and what we ought to do in multiple option cases might differ significantly. Limited Aggregation views in multiple option cases face a number of extra complications. Namely intransitivity and the *principle of irrelevant alternatives*.⁵⁶ Plus the related problem of deciding how we ought to determine relevance when claims compete across multiple groups.

I do not see these issues as insurmountable, and I believe the account I provide here might provide a solid basis for such an extension. For instance, on the latter issue, my thesis already deals with competing requirements of relevance, and the question of how claims can be relevant to different strength claims across heterogeneous cases. How we deal with competing requirements of fairness across multiple option cases, is likely to rely on some of the same ideas and thus extending the account should be achievable. However, we will need a different, and potentially much more complex, sequence to capture what we ought to do in such cases. At the very least, some of the insights, principles and structures uncovered in this thesis should assist in navigating a way through those difficulties.

⁵⁶ See Norcross (2002) and Halstead (2016) for criticisms of aggregative views based on intransitivity. See Temkin (2012) and Rachels (1998) for defences of the general intransitivity of better than. See Kamm (2007, pp. 484-486) and Voorhoeve (2013) for defences of intransitivity in aggregation cases.

Conclusion

Secondly, three related issues that I have not been able to cover are risk, uncertainty, and intrapersonal aggregation.⁵⁷ I briefly touch on risk (and even more briefly intrapersonal aggregation) in Appendix 1 to discuss a second dilemma that Horton poses for limited aggregation views.

However, developing a positive account of what we ought to do in such cases falls out of the scope of this thesis. Nevertheless, understanding how my account, and limited aggregation views more generally, accommodate such concerns will be vital in developing a complete, robust, and adaptable account of what we ought to do in such situations. I hope to return to these questions in future work, alongside the multiple option cases.

It also goes without saying that this thesis only covers some of our moral reasons. Questions of aggregation, and in particular resource distribution, cannot be considered in complete isolation from these wider moral considerations. There may be domain specific moral reasons that alter how we aggregate and which features are relevant or irrelevant to our decision making. For instance, in healthcare we might deem economic benefits irrelevant to the decision. There will also be considerations of justice and rights (including those of groups) which might differ from the considerations we have already accounted for that need to be incorporated. Nevertheless, the account I develop here should, I hope, provide a basic structure and ground for considering these wider features.

Lastly, I would like to touch on some of the upshots and takeaways from the thesis. Limited aggregation has many potential applications, especially when we consider questions of resource allocation. I mention a number of these applications in the introduction to the thesis so I will not list them again here. Of course, how insights from my thesis affect how we make decisions in each of those domains will vary. Whether and how my sequential account can be applied in practice remains an open question, but one I plan to return to.

⁵⁷ For some recent work on risk see Horton's (2020) critique and Patrick Wu's (2022) defence. For work on limited aggregation and uncertainty see Steuwer (2021b).

Conclusion

I also think there are some methodological (and decision procedural) benefits to my sequential approach worth mentioning. An epistemological issue complex accounts of aggregation (and morality more generally) face is that there are often too many features for us to accurately have intuitions about or to be able to correctly judge. For instance, judgements about complex heterogenous cases with many different claims are too difficult to judge without principles or a guide to how we ought to weigh our reasons.

The benefit of my account is that it breaks these difficult cases down into more manageable chunks, over which moral judgement is possible. We only need to make moral judgements about the relevance of each claim compared to another claim when those claims come into conflict, and when they do come into conflict, we can avoid all the other features whilst we determine how these claims interact. We can then build up a picture out of these individual judgements that better captures everything we want to say about these cases. In this way, the account can bring together disparate intuitions and combine them in such a way to justify those intuitions and explain the moral phenomena better.

It strikes me that this approach would thus be highly valuable, not just in cases of limited aggregation, but more widely in moral philosophy and applied ethics. In so far as this thesis develops a plausible sequence, I hope it provides further impetus to develop more such accounts and explore the feasibility of these style approaches more generally.

Furthermore, this kind of account is also potentially amenable with technology. Sequential approaches may prove difficult for human decision-makers to use accurately due to their complexity. This is especially true when decisions need to be made under time pressures. However, computers have no such trouble implementing sequences (or algorithms) and thus might be a fertile ground for the use of these sequences. For instance, self-driving cars or other automated machines could be programmed with these sorts of sequential decision-making systems (assuming they were

Conclusion

also provided with the right values) to help them make, non-utilitarian, moral decisions in a wider variety of cases, including where contextual interactions might alter what we ought to do.

Obviously, the extent to which the sequential approach might help guide ethical decision-making in technology is way outside the scope of this thesis. Nevertheless, I think this is a potentially important takeaway that ought to be given more consideration. Much has been said recently about the ethical problems algorithms pose, but little has been said on how algorithms (or sequences) might help solve ethical problems. My contention is that by focusing more on the second question we might also find answers to the first.

Appendices

Appendix 1 – Risk, Reductio and Weak Additivity

In Chapter 1 I demonstrate why Joe Horton's 'fatal dilemma' for limited aggregation views fails, and why it relies on the faulty assumption of *Weak Additivity*. However, Horton also provides a second critique of limited aggregation views. This critique focuses on cases involving risk, and supposedly shows a reductio for limited aggregation views. For reasons of space, my thesis does not deal directly with cases involving risk, though I hope to expand on this elsewhere. However, it is worth briefly noting Horton's reductio, and where I believe the reductio goes wrong.

Firstly, Horton's reductio. He asks us to consider the three following cases:

Villain 1: A villain has kidnapped A and B. He will either (1) inflict a migraine on A, or (2) inflict a one-in-a-zillion chance of death on B. You must choose which.

Villain 2: A villain has kidnapped ten zillion X people and ten zillion Y people. He will either (1) inflict a migraine on each X person, or (2) randomly select and kill ten Y people. You must choose which.

Villain 3: A villain has kidnapped ten zillion X people and ten zillion Y people. He pairs each X person with a Y person. For each pair, the villain will either (1) inflict a migraine on the X person, or (2) give the Y person a ticket for a lottery with ten zillion tickets. You must choose between these options for each pair in turn. You know that, after you have chosen for each pair, the villain will randomly select ten tickets and kill anyone who has a corresponding ticket. (Horton, 2020, pp. 516 - 517)

Horton argues that limited aggregation views ought to choose option (2) in *Villain 1*, that they are compelled to choose option (1) in *Villain 2*, and that *Villain 3* is in "all morally relevant respects

equivalent” (Horton, 2020, p. 517) to both *Villain 1* and *Villain 2*. Thus, there is an inconsistency, whichever option we choose in *Villain 3*.

The first thing to note is that I am unsure whether the limited aggregationist is committed to choosing (2) in *Villain 1*. Horton himself notes that we “could argue that the relevance of competing claims is determined not by their relative strength, but rather by the relative size of the burdens that ground them” (Horton, 2020, p. 525). If this is the case, then the migraine is not relevant to even the very small chance of death, and so we ought to choose option (1).

Horton argues that this is implausible for a couple of reasons. Firstly, he compares *Villain 1* with day-to-day cases in which we take minor risks of major harms in order to reduce minor burdens. For instance, we think it is permissible to drive to the pharmacy to pick up migraine medication, even though doing so risks causing death via road accident.

However, I think this comparison is an unhelpful disanalogy for several reasons. I will list these reasons briefly rather than explain them in detail. Firstly, the risk in the pharmacy case is foreseen but not an intended or necessary part of the choice, whereas in *Villain 1* it is. Secondly, the risk in the pharmacy case is not a novel risk; the risk is a background risk inherent in all day-to-day activities. In so far as we all agree to operate in a society with such risk, we have consented to that risk.

Thirdly, actions like driving to the pharmacy should not be considered in isolation; doing so will mislead our intuitions. Driving to the pharmacy is a part of a wider series (or set) of actions that together make up different more valuable wholes, such as ‘the freedom to drive’, ‘living a life without overwhelming restrictions’ and even ‘living a good life’. Whilst each individual trade-off might not be worth the risk, the sum of the trade-offs, as a whole, is. This has been particularly clear post-covid: each individual day in lockdown was a minor burden, but collectively the entire lockdown was a major burden. This was not simply a sum of the minor burdens, but a qualitative difference as the isolation compounded.

Horton's second argument is that "it is better to be subjected to a tiny chance of death than to suffer a migraine. That is why it is rational to take aspirin, despite the tiny chance of suffering a fatal allergic reaction. So, a claim to be saved from a migraine is stronger than a claim to be saved from a tiny chance of death" (Horton, 2020, p. 524). Horton argues that this means A has a stronger claim than B in *Villain 1*.

However, there are two reasons why this does not follow. Firstly, the case of taking aspirin is an example of intrapersonal aggregation not interpersonal aggregation. What it is rational and permissible to prefer in cases of intrapersonal aggregation might differ significantly from interpersonal aggregation. It is permissible to risk your life driving at high speeds on the Nürburgring, but it is not permissible to risk others' lives driving at high speeds down country lanes.

Furthermore, it does not follow from one option being preferred that that option generates a stronger claim. One might legitimately prefer the risk, but that does not necessarily make the riskier option the stronger claim. We see this in other similar cases: I might prefer that the hospital gives me £1000 rather than spend £2000 on a treatment for a broken arm, but I only have a claim to the latter.

Thus, if the limited aggregationist is forced to bite a bullet in one of these cases, it would be fairly straightforward to bite the bullet in *Villain 1*. Limited aggregation views are premised on giving significant priority to those with more at stake, considering their personal perspectives or partial concern, and giving each individual a justification for your action. It strikes me as permissible but not required for somebody to take a very small risk of dying to cure the migraine of another, and that without their consent we ought to choose option (1).

Nevertheless, with all of the above said, one might still be inclined to agree with Horton. Suppose that we ought to choose option (2) in *Villain 1*. In this case, I still think Horton's reductio fails, and it fails for essentially the same reason that his 'fatal dilemma' fails. Namely, it assumes *Weak Additivity* once again, although less obviously:

Appendix 1 – Risk, Reductio and Weak Additivity

Weak Additivity: If A is preferable to B, and C is preferable to D, then A with C is preferable to B with D.

Horton's reductio relies on this principle to say that *Villain 1* and *Villain 3* are morally equivalent. *Villain 3* consists of ten-zillion *Villain 1*s: we have *Villain 1₁*, *Villain 1₂* all the way to *Villain 1_{Ten-Zillion}*. If option (2) is preferable to option (1) in *Villain 1₁*, and option (2) is preferable to option (1) in *Villain 1₂* (and so on), then option (2) in *Villain 1₁* with option (2) in *Villain 1₂* (and so on) is preferable to option (1) in *Villain 1₁* with option (1) in *Villain 1₂* (and so on).

This seems fairly natural and explains the force of the reductio. However, as we see in Chapter 1, *Weak Additivity* does not always hold and should not be assumed. What might be acceptable at the individual level does not always hold at the general level. However, if we are going to make an appeal to the failure of *Weak Additivity* here, we need to explain why each individual choice is not independent of the other choices in *Villain 3* and what interaction effects there are.

It might seem like these choices are completely independent of each other, but it is important to see the role the lottery plays in *Villain 3*. There are ten-zillion tickets of which ten will lead to death. Now, whichever way we formulate the lottery, each independent choice affects the pool of tickets, and the chance of another individual dying.

Let us first suppose that the lottery is run like a golden ticket lottery. The death tickets are hidden amongst all the non-death tickets. In each individual choice, if we pick a ticket we check if it is a death ticket before moving onto the next choice. If the lottery is set up like this, then each individual who picks a non-death ticket reduces the chances that the next person will pick a non-death ticket. The knock-on effect might be very minor, but it is sufficient such that each choice is not independent of each other. Furthermore, the knock-on effect might only be minor to each other individual choice, but on aggregate this knock-on effect is significant. For instance, a 0.00001% increased chance of death to each of 100,000,000 people equates to one extra expected death in that group.

Appendices

Suppose instead that the lottery is run as a raffle. Only after each individual decision is made is the raffle run and the death tickets chosen. If the lottery is set up like this, then in each individual choice, every choice of a raffle ticket reduces the chances that the previous person will end up with a death ticket. However, the first individual choice, at the time we make the decision, is the only raffle ticket in the lottery and so (if no other individuals join the raffle) is guaranteed to pull a death ticket. We might think that this does not permit choosing the raffle ticket option in the first individual choice. If it is permissible to choose the raffle ticket option in the first individual choice, it is only so because the raffle ticket option will also be chosen in the other individual choices. Again, this shows the interdependence of these decisions.

Lastly, suppose that the lottery is run with ten-zillion tickets and ten tickets will be chosen at random after each individual decision has been made. This seems to be the way that Horton develops the lottery in *Villain 3* above. If the ticket that is chosen matches somebody, then that person dies; if not, then nobody dies. In this case, it might seem like the choices are independent. Each individual choice does not affect the chances that somebody else dies: if the ticket on offer is pulled, then the person in this case dies, and if not, nobody dies. However, whilst the chances of other people dying does not change, it does not mean the choices are completely independent.

I think there are two things to note here. Firstly, this case looks in some ways very similar to the golden ticket style lottery above. The only difference is with the time at which we find out who has a death ticket and who does not. We might think this aspect should not make any difference. From the post-hoc perspective, we can see that each individual choice of a ticket which is not eventually pulled affects the chance of the next person picking a ticket that will become the death ticket. Perhaps, there is something dissimilar between this way of running a lottery and the golden ticket way. But if such distinctions matter, then I cannot see how Horton could maintain that *Villain 1*, *Villain 2* and *Villain 3* are morally equivalent, as the differences in the cases are equally subtle.

Secondly, even if we maintain the ante-hoc perspective, this way of running a lottery still means that some decisions are affected by prior decisions. For instance, the last individual choice is different to the first individual choice, because of the choices that come before it. The last individual choice (in fact if there are ten death tickets, the tenth-last individual choice), guarantees that someone will die. With each individual decision, the chance of somebody dying is increased by a very small amount. However, the last decision shifts the choice from almost certain that someone will die, to absolutely certain that someone will die. Similarly, the middle choice shifts the balance of probability from more likely that no one will die, to more likely that someone will die. Again, this shows a mild interaction effect between each choice. Now these effects are very minor so we might be tempted to dismiss them, but that would be an error. Ironically, these minor effects aggregate such that they matter and these mild interaction effects on the whole are sufficiently large to prevent us from being able to assume *Weak Additivity*.

Once we see these interaction effects, I think it is easy to see why it might be permissible to choose option (2) in *Villain 1* but why we must choose option (1) in *Villain 2* and *Villain 3*.

Now let us consider different versions of *Villain 2* and *Villain 3* that do not rely on lotteries:

*Villain 2**: A villain has kidnapped ten zillion X people and ten zillion Y people. He will either (1) inflict a migraine on each X person, or (2) inflict a one-in-a-zillion chance of death on each Y person. You must choose which.

*Villain 3**: A villain has kidnapped ten zillion X people and ten zillion Y people. He pairs each X person with a Y person. For each pair, the villain will either (1) inflict a migraine on the X person, or (2) inflict a one-in-a-zillion chance of death on the Y person. You must choose between these options for each pair in turn. (Horton, 2020, pp. 518-519)

In *Villain 3**, each individual decision does seem independent. Thus, there are no interaction effects between the individual decisions. Thus, *Weak Additivity* can be assumed in this case. As such, in

Appendices

*Villain 3** it seems permissible to me to choose option (2) in each case, even though it is an almost statistical certainty that at least one person will die. However, *Villain 3** is not morally equivalent to *Villain 2*. In *Villain 2*, like *Villain 3*, it is an absolute certainty that someone will die, and the migraines can only be cured if people die.

Nor is *Villain 3** perfectly identical to *Villain 2**. The move from *Villain 3** to *Villain 2** once again requires *Weak Additivity*, and there are some interaction effects in the move. In *Villain 3** the salvation of each migraine only comes into competition with the tiny chance of death of one other person. Whilst in *Villain 2** each migraine comes into competition with the tiny chance of death of a very large number of people. Similarly, in *Villain 3** each individual tiny chance of death only comes into competition with one migraine, but in *Villain 2** each individual tiny chance of death comes into competition with many migraines. It is plausible that these interaction effects mean that *Villain 3** and *Villain 2** should be treated differently, and I think a limited aggregationist is well within their right to choose option (2) in each *Villain 3** case but choose option (1) in *Villain 2**.

To see why, consider the personal perspective of the claimants in these cases. If I was a migraine claimant in *Villain 3**, I might legitimately prefer my own migraine to be cured at the risk of one other person having a tiny chance of dying. But in *Villain 2** we might think it is impermissible for me to prefer my own migraine to be cured rather than ten-zillion people all having that same tiny chance of dying. Or we might say that I can legitimately prefer my own migraine to be cured at the very small statistical chance of someone dying in *Villain 3**, but that I cannot prefer my migraine to be cured at the very large statistical chance of at least one person dying in *Villain 2**.

Nevertheless, one might maintain that these interaction effects either make no difference or cancel each other out, and that we ought to treat *Villain 3** and *Villain 2** the same. If so, this would mean choosing option (2) in *Villain 2** and we might think this moves significantly away from the limited aggregation perspective. However, even if this is the case, *Villain 2** and *Villain 2* also strike me as distinct cases. Distinct enough that a limited aggregationist could plausibly treat them differently. In

Villain 2 ten people are guaranteed to die, whilst in *Villain 2** the deaths are only a statistical certainty. This might seem like a minor difference, but I think this difference can be important.

Consider a case where a patient may undergo surgery in an attempt to regain the ability to walk. However, this surgery has a 1% chance of ending in the patient's death. We think that it is permissible for the patient to undergo the surgery. Now, if 10,000 people undergo the surgery, statistically 100 of them will die. However, we also think it remains permissible for the 10,000 patients to take this risk, and for surgeons to undertake the surgery knowing that some of their patients will likely die.

However, consider another case, where we have 10,000 people who cannot walk. To give them the ability to walk we need a donor spinal cord. One person's spinal cord would be sufficient to give ninety-nine others the surgery needed to regain the ability to walk. But there are no donors, and the patients are running out of time. Thus, to give each patient a chance to walk again we could run a lottery. In this lottery each individual has a 1% chance of being the donor and 99% chance of regaining their ability to walk. Of course, the donor will die as their spinal cord is extracted. In this case, it strikes me as impermissible to run this lottery. The risk to each individual patient is identical and statistically the same number of patients will die, but there is a significant moral difference in these two cases.

In the first case, the deaths are statistical. Each individual trade-off is independent, and any deaths are unlucky and not a necessary part of the choice. However, in the second case the deaths are not statistical. The ability of one person to walk strictly requires that another person dies. The deaths in this case are not unlucky and merely foreseen; they are necessary part of the choice. The difference between *Villain 2* and *Villain 2** is the same. In *Villain 2*, the deaths are foreseen but individually unlucky and not necessary for the curing of the migraines. In *Villain 2** the deaths are required for the migraines to be cured.

Lastly, these issues surrounding risk are not unique to limited aggregation views. Rational choice under conditions of risk affects all views, not least fully aggregative views. Firstly, consider another villain case where we alter where the risk lies:

Villain 4: A villain has kidnapped ten million zillion X people and A. He will either (1) kill A, or (2) inflict a one-in-a-zillion chance of a migraine on each X person. You must choose which.

Of course, we can strengthen this even further:

*Villain 4**: A villain has kidnapped ten million zillion zillion X people and a zillion Y people. He will either (1) kill the Y people, or (2) inflict a one-in-a-zillion chance of a migraine on each X person. You must choose which.

Surely, in both these cases we ought to save the people from dying and choose option (2). Of course, fully aggregative views will say that statistically the number of migraines that will be cured will outweigh the lives, and so we should choose option (1). This is not a formal dilemma for the fully aggregative view, and it does not attempt to demonstrate an inconsistency in the fully aggregative position. However, this consequence nevertheless seems significantly worse than the supposed inconsistency that limited aggregation views face. If the introduction of risk has no effect on permissibility, such that *Villain 1* and *Villain 4* ought to be treated the same, then *Villain 1* strikes me as the clearly the right bullet to bite.

Furthermore, these problems regarding risk share the same underlying structure as a large number of other problems with dynamic choice.⁵⁸ In particular, I think it is interesting to note that underlying these cases is a sorites paradox. Each step in a sequence makes no moral difference but the sequence as a whole does make a difference. Pulling one of your hairs out as you sleep (i.e.,

⁵⁸ In particular, see Patrick Wu's (2022) discussion of Horton's reductio. Wu demonstrates that Horton's reductio can be simplified to remove risk entirely and focus on sequences of action. He then identifies three diachronic principles that Horton's reductio relies on and demonstrates why we might find each of these principles implausible. Whilst Wu does not mention it himself, it is clear that all these principles fail because they assume *Weak Additivity* and ignore important interaction effects.

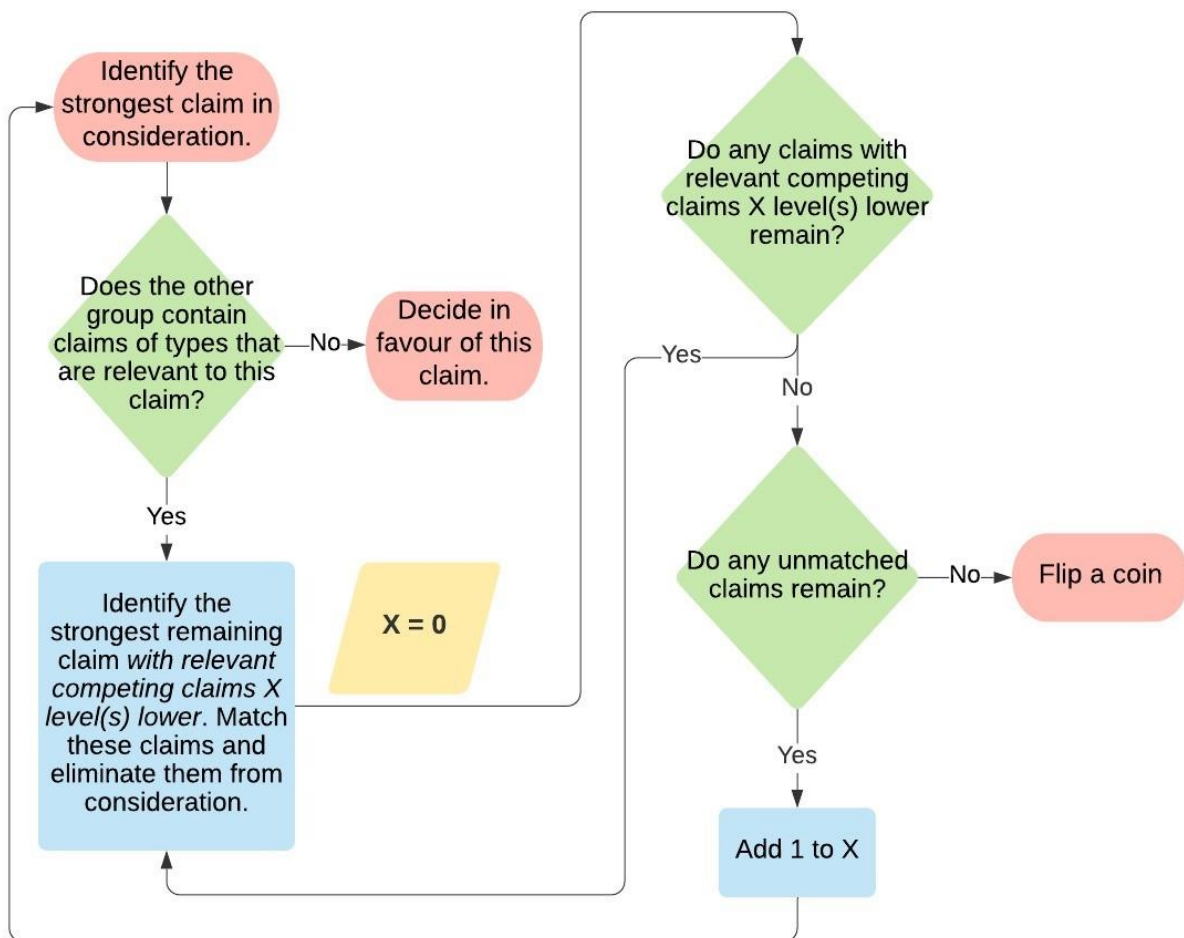
painlessly) will not make you bald and is thus permissible. Nor will the second, or third etc. However, pulling out all of your hair as you sleep (again painlessly) will make you bald and so is impermissible.

It is not clear how we ought to handle such cases. But it is clear that such problems are not unique to limited aggregationists. That limited aggregationists face this problem in the *Villain* cases, only demonstrates that limited aggregationists face one extra particular instance of a general problem. This extra instance should not count against such views and rejecting limited aggregation views on these grounds strikes me as is inappropriate.

To conclude my brief foray into risk, there are a whole host of different plausible options for the limited aggregationist to take in these cases, and limited aggregationists might reasonably disagree about the best approach. Horton has nevertheless helpfully shone a light on another ambiguity that limited aggregationists might need to address if they are to have a complete view. Note that limited aggregation views are hardly the only views that need to address this problem to have a complete view. Sadly, my thesis is not the place to address these important questions. Hence, I shall leave it open-ended as to which of the options limited aggregationists ought to take to avoid Horton's reductio.

Appendix 2 – Sequential Claims-Matching by Closeness – Shortened

- I. Identify the strongest claim in consideration. Does the other group contain claims of types that are relevant (i.e., sufficiently strong relative to) to this claim and sufficient in number to match it?
 - a. If so, go to Step II.
 - b. If not, you should decide in favour of the group with the strongest claim in consideration.
- II. Identify the strongest remaining claim *with* relevant competing claims [zero] levels lower. Match these claims and eliminate them from consideration.
- III. Continue step II until there are no claims with relevant competing claims [zero] levels lower.
- IV. Repeat steps I - III, with competing claims one level lower, then two levels lower, and so on.
- V. Continue until either:
 - 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).



Appendix 3 – Fairness Last

In Section 3 of Chapter 6, I develop my complete account of aggregation by addressing S-claims first. Only after all S-claims are considered, and those S-claims that can be taken out of consideration are, do we consider the T-claims and F-claims. I shall call this approach to the order of claims-matching *Fairness Last*.

However, I do not provide an argument for the *Fairness Last* approach, and we could match S-claims and F-claims in a different order. So, what order should we match S-claims and F-claims in? Should we match all S-claims or F-claims first? Or should the order that we consider the S-claims and F-claims depend on other features?

As far as I can see, there are a few potential approaches to the order of matching. These three seem to me to be the strongest contenders:

Fairness Last: Match all S-claims with each other first.

Fairness First: Match F-claims at the earliest opportunity.

Closest First: Give neither type of claim priority. Match claims by closeness of the underlying stake. I.e., match a death claimant's S-claim with another death claimant's F-claim before matching the death claimant's S-claim to a paraplegia claimant's S-claim.

All three options seem to have some initial plausibility. *Fairness Last* is the simplest option: F-claims only come into play when all S-claims have been considered. *Fairness Last* also seems the most respectful option - we might think it is more respectful to match S-claims with other S-claims before we match them with any F-claims.

Fairness First is a little more complicated: F-claims would need to be matched by S-claims every time F-claims arise. Nevertheless, *Fairness First* has some initial plausibility too. It would mean that we consider the requirements of fairness as we go, rather than all at once at the end. If one thinks that fairness concerns have some sort of priority, then this approach would also capture this.

Appendices

Closest First is the most complicated: instead of giving priority to either type of claim, it compares claims by their strength. For instance, a lost limb claimant's F-claim is closer in strength to somebody else's lost limb S-claim, than a dying person's S-claim is to that same lost limb S-claim. Thus, it might be more appropriate, or respectful, to match a lost limb claimant's S-claim to another lost limb claimant's F-claim, rather than to the death claimant's S-claim. In this way *Closest First* seems closest to *Match by Closeness*.

However, whilst *Fairness First* and *Closest First* have some initial plausibility, both options share a fatal flaw. By taking into consideration F-claims before all S-claims are matched, both options allow F-claims to form part of the reason to save one group outright. This is deeply problematic as F-claims are only claims to have a fair chance of being saved, and thus they cannot be part of a reason not to give others a fair chance to be saved. To see why consider Case 16:

Case 16

Level	Group A	Group B
1 - Death	4 <4>	4 <4>
2 - Paraplegia	1	
3 - Lost Limb		
4 - Lost Hand		10

First let us consider *Fairness First*. In this case, both Group A and B have death claimants with F-claims, as such we must ask which F-claims we should match first? Do we match the F-claims in Group A or in Group B first? It is not immediately clear what the right option is, and so this approach brings with it a further ambiguity. However, this is not the main problem.

Suppose we decide to match the F-claims in Group B first. For arguments sake, let us also suppose that a paraplegia claimant's S-claim matches with exactly four death claimants' F-claims. As such, the paraplegia S-claim and the death claimants' F-claims are taken out of consideration. We then must consider the lost hand claimants' S-claims. Suppose these S-claims match and outweigh the death claimants' F-claims in Group A, and thus we ought to save Group B outright.

Appendix 3 – Fairness Last

We can see then that the F-claims in Group B, by balancing with the paraplegia claimant's S-claim in Group A, has enabled the lost hand claimants' S-claims to be decisive in favour of Group B. The F-claim has formed part of the reason why we ought to save Group B outright, even though F-claims ought only to support the use of a fair lottery.

This makes *Fairness First* untenable, but it also makes *Closest First* untenable as *Closest First* will give us exactly the same answer in Case 16. In this example a paraplegia claimant's S-claim is closer to a death claimant's F-claim than it is to a lost hand claimant's S-claim. Thus, *Closest First* will match the F-claims and S-claims in the exact same manner as *Fairness First*. So once again, the F-claims in Group B, which are only reasons to give Group B a fair chance to be saved, have formed part of the reason to save them outright.

Fairness Last on the other hand, does not encounter this problem. All S-claims will be considered and matched with other S-claims before any F-claims are considered. Thus, F-claims will only ever prevent S-claims from being decisive rather than form part of a reason to save one group outright. Only *Fairness Last* can capture this important feature and so it is vital that the account I develop matches all the S-claims first.

This problem is also why it is helpful to distinguish between S-claims and T-claims. F-claims cannot outweigh just any S-claim, they can only outweigh S-claims that are T-claims, i.e., S-claims that cannot be outweighed by other S-claims.

Appendix 4 – Match in Order of Strength

Appendix 3 addresses the order in which we ought to match S-claims and F-claims; however, it only solves part of the issue regarding the order of claims matching. Remember from Chapter 2 that there are two aspects regarding matching claims – the order we match claims, and which claims we match each claim to. *Fairness Last* answers the first of these aspects; however, we still need to know which F-claims ought to be matched to which T-claims. To see why this matters, consider Case 17:

Case 17

Level	Group A	Group B
1 - Death	10 <10>	10 <10>
2 - Paraplegia		
3 - Lost Limb		1
4 - Lost Hand	1000 <1000>	1000 <1000>
5 - Lost Finger		1000

In this case we have two groups of F-claims of differing strengths and two groups of T-claims of differing strength. How we match these F-claims and T-claims will affect whether we save Group B outright or use a lottery.

Firstly, suppose we match the death claimants' F-claims with the lost limb claimant's T-claim, taking the lost limb claimant's T-claim out of consideration. The lost hand claimants' F-claims are thus only needed to outweigh the lost finger claimants' T-claims - which they do - and so we ought to run a lottery and give everyone an opportunity to be saved.

Conversely, suppose we match the death claimants' F-claims with the lost finger claimants' T-claims, such that both sets of claims are taken out of consideration. Now, suppose the lost hand claimants' F-claims are not relevant to the lost limb claimant's T-claim. If so, no matter the number of lost hand claimants' F-claims, the lost limb claimant's T-claim will be decisive in favour of Group B.

Thus, which T-claims are matched to which F-claims makes a difference to who we ought to save in heterogenous cases. So how should we match the F-claims and the T-claims in these cases. What is the matching rule we ought to use? Well, let us consider three different matching rules:

Appendix 4 – Match in Order of Strength

Strongest Decides: T-claims should be matched in whatever way is in the interest of strongest overall claimant.

Match in Order of Strength: T-claims should be matched in order of strength. I.e., we match the strongest T-claim first and then the second strongest T-claim, and so on.

Match by Closeness: when matching T-claims, where possible match T-claims with closest relevant F-claims (where closeness is determined by the underlying stake for each claimant).

How does each of these principles handle Case 17?

Firstly, *Strongest Decides* would not handle it very well, since there is no strongest overall claimant, and so the order in which we ought to match the claims would remain ambiguous. As such, I think we can rule *Strongest Decides* out. Perhaps a more precise version of *Strongest Decides* could be adopted. Call it *Strongest Tiebreaker Decides*:

Strongest Tiebreaker Decides: T-claims should be matched in whatever way is in the interest of strongest tiebreaking claimant.

This principle is not ambiguous, and there is a certain logic to it - *Strongest Tiebreaker Decides* would favour the person with the most to lose who does not have their S-claim matched by another's S-claim. However, I think this principle goes too far. *Strongest Tiebreaker Decides*, much like the version of *Strongest Decides* we see in Chapter 2, seems to give too much influence to the strongest T-claim. It allows the claimant with the strongest T-claim to gerrymander how we match claims, rather than provide a strong independent principle which all claimants could agree to in advance.

Strongest Tiebreaker Decides always favours the tiebreaking group, such that if there is any way to compare claims such that the T-claims win, this would be the way to match claims. This implication would, of course, depart fairly significantly from the reasoning given for *Strongest Decides* by Van Gils and Tomlin (Van Gils & Tomlin, 2019, pp. 250-252) (also see Chapter 2 of this thesis) as we are no longer giving priority to those with the most at stake but to those with less (perhaps much less) at

stake but who have not had their S-claims matched by competing S-claims. In some cases, *Strongest Tiebreaker Decides* will require us to treat weak (but relevant) T-claims as having priority over much stronger F-claims.

Match in Order of Strength on the other hand seems less arbitrary. In Case 17, *Match in Order of Strength* tells us to match the lost limb claimant's T-claim first and to match this T-claim to the strongest (relevant) competing F-claim – in this case the death claimants' F-claims. Supposing these F-claims match exactly with the lost limb claimant's T-claim, *Match in Order of Strength* then tells us to compare the lost finger claimant's T-claims with the strongest remaining F-claims (i.e., those of the lost hand claimants). On the assumption that the remaining F-claims outweigh the lost finger claimants' T-claims, then *Match in Order of Strength* tells us to run a lottery. This seems like an appropriate answer in this case.

Match by Closeness also tells us to run a lottery. However, it might be a little more difficult to see why. It seems like the lost limb claimant's T-claim and the lost finger claimants' T-claims are equally close to the lost hand claimants' F-claims. Thus, it might seem that there is an ambiguity as to how we match the T-claims.

However, the reality is that this ambiguity does not arise because of the *Strongest Always Tiebreaks* principle outlined in Appendix 7. *Strongest Always Tiebreaks* states that F-claims are irrelevant to stronger T-claims. This means T-claims can only be matched and outweighed by F-claims with as strong, or stronger, underlying stake, i.e., F-claims on the same level or higher.

Thus, in Case 17, the lost finger claimants' T-claims are the only T-claims with relevant F-claims one-level apart. So, according to *Match by Closeness*, we compare these T-claims with the lost hand claimants' F-claims, and then compare the lost limb claimant's T-claim with the death claimants' F-claims. This results in all the T-claims being matched and outweighed, and thus *Match by Closeness* also tells us to run a lottery. So, *Match by Closeness* and *Match in Order of Strength* get the same answer in this case.

Appendix 4 – Match in Order of Strength

This is no coincidence. In fact, *Match by Closeness* and *Match in Order of Strength* will get the same answer in all cases, so long as we maintain *Stronger Always Tiebreaks* and *Anchoring Relevance* as per Appendices 6 & 7. If F-claims are never relevant to stronger T-claims, and T-claims are only relevant if they are globally relevant to the strongest competing S-claim, then there are no cases where these two matching principles will give different answers. The principle we choose might alter the order in which we match the claims, but the answer will always be the same.

To see why *Match by Closeness* and *Match in Order of Strength* will always get the same answer when we maintain *Anchoring Relevance* and *Stronger Always Tiebreaks*, we should consider the cases in which *Match by Closeness* and *Match in Order of Strength* differ in the normal cases. As we saw in Part One of the thesis, *Match by Closeness* and *Match in Order of Strength* differ in cases like the First Horn Case:

First Horn Case

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia		
3 - Lost Limb	(100)	(6)
4 - Lost Hand		1000

In this case, there is an ambiguity about how we compare and match the various claims. We could match the A1 claim with the B3 claim and the A3 claim with the B4 claim, or we could match the A3 and B3 claims together and compare the A1 claim to the B4 claims. How we settle this ambiguity determines which group we save.

As we saw previously, *Match in Order of Strength* tells us to match the A1 claim with the B3 claims. Then the A3 claims with the B4 claims, such that Group B wins. *Match by Closeness* on the other hand tells us to match the A3 and B3 claims, such that the B4 claims must compete with the A1 claim. The B4 claims are irrelevant to the A1 claim, so *Match by Closeness* tells us to save Group-A. Thus, the difference between the two groups is how they decide to match the strongest S-claim in Group-B.

Appendices

This ambiguity arises because we treat S-claims with local relevance. This means that the B4 claims are relevant to the A3 claims but not the A1 claim. Thus, the way in which we compare the claims makes a big difference as to whom we save. An important element of this is that S-claims can be compared and matched with both weaker and stronger S-claims.

Now let us reconsider Case 17 as an example of a tiebreaking case:

Case 17

Level	Group A	Group B
1 - Death	10 <10>	10 <10>
2 - Paraplegia		
3 - Lost Limb		1
4 - Lost Hand	1000 <1000>	1000 <1000>
5 - Lost Finger		1000

In Case 17, and all other tiebreaking cases, T-claims are considered relevant or irrelevant according to global (or anchoring) relevance, i.e., T-claims are only relevant if they are relevant to the strongest overall F-claim. Thus, there is no way of matching T-claims that makes them relevant in one matching and irrelevant in another matching. Thus, there is no ambiguity about whether a T-claim is relevant or not.

However, F-claims are considered relevant according to local relevance. Thus, if we compare the level 4 F-claims with the level 3 T-claim, they are irrelevant. But if we compare the level 4 F-claims with the level 5 T-claims, they are relevant. So, the way in which we match the F-claims will still make a difference as to whether those claims are relevant.

Nevertheless, *Match by Closeness* and *Match in Order of Strength* will always get the same result in these cases. Again, this is because of the *Strongest Always Tiebreaks*. *Strongest Always Tiebreaks* states that T-claims can only be matched by as strong, or stronger, F-claims, but not to weaker F-claims. This means that there is no ambiguity about whether to match T-claims with stronger or weaker F-claims. Thus, the only remaining ambiguity is which T-claims do we match to which stronger F-claims.

Appendix 4 – Match in Order of Strength

Whether we start from the strongest T-claim and match it to the strongest competing T-claim, or whether we start with the T-claim with the closest competing F-claim, no longer matters. Because of the above limitations both approaches will lead to the same T-claims and F-claims being relevant, and thus the same answer.

So, it seems that the principles agree, we ought to run a lottery in Case 17. I think this is the right answer. I struggle to see why we would prefer to match the lost finger claimants' T-claims to the death claimants' F-claims, and thus leave the lost limb claimant's T-claim to trump the lost hand claimants' F-claims and be decisive in favour of Group B. The only argument I can see for why we should match the claims in such a way as to save Group B outright would be to say that it is more appropriate to match the death claimants' F-claims with the lost finger claimants' T-claims than with the lost limb claimant's T-claim. But what would be the reasoning here?

Perhaps we could say that the lost finger claimants' T-claims are closer in strength to the death claimants' F-claims than the lost limb claimant's T-claim is because F-claims are significantly weaker than S-claim. This might mean that a dying person's F-claim has weight closer to the weight of a lost finger claimant's T-claim than to a lost limb claimant's T-claim. Following the logic behind *Match by Closeness*, we ought then to match the death claimants' F-claims with the lost finger claimants' T-claims rather than the paraplegia claimant's T-claim. For lack of a better term, I shall call this *Weight Adjusted Match by Closeness*:

Weight Adjusted Match by Closeness: when matching T-claims, where possible match T-claims with closest relevant F-claims (where closeness is determined by the weight of the F-claim).

However, whilst there is some logic behind *Weight Adjusted Match by Closeness*, I think it takes the logic of *Match by Closeness* in the wrong direction. Limited aggregation views are concerned with what is at stake for each individual from their perspective, and *Match by Closeness* is an attempt to

Appendices

capture this. Matching claims by closeness (of underlying stake) ensures that each claimant has what is at stake for them compared to something similarly at stake for someone else.

Conversely, if we compare claims by the closeness of the weight or strength of their T-claim or F-claim rather than with what is at stake for them, we step away from their personal perspective. The claimants are not concerned with the weight of their claims but with what they have to lose if they are not saved. In this way, if we match claims by the closeness of the strength of claims I think we move too far away from the initial reasoning that got us to *Match by Closeness*.

Furthermore, whom would *Weight Adjusted Match by Closeness* be more respectful of? Definitely not the strongest claimants. For instance, in Case 17 above it would, obviously, be more reasonable to explain to the death claimants that they cannot have a chance to be saved because there are people who could be saved from lost limbs rather than because there are people who could be saved from lost fingers. Nor do I think it is any more respectful of the weaker claimants; lost finger claimants are not better respected by comparing their T-claims to death claimants' F-claims rather than lost hand claimants' F-claims.

If we want to give those with the most at stake any sort of priority, as limited aggregationists do, then we should want to prioritise matching the strongest claimants' claims (in this case F-claims) in the most respectful way possible. *Weight Adjusted Match by Closeness* (and *Strongest Tiebreaker Decides*) do the exact opposite in these cases by matching the strongest claimants' F-claims with the weakest relevant T-claims rather than the strongest relevant T-claims.

With this in view then, I think we can safely proceed to settle the ambiguity in Case A above, and other similar cases, in favour of running a lottery. We ought to match the strongest T-claims with the strongest or closest relevant F-claims. Whether we use *Match by Closeness* or *Match in Order of Strength* does not matter too much as we get the same answer either way.

Appendix 4 – Match in Order of Strength

Nevertheless, I use *Match by Closeness* for the complete account in Chapter 6. This is really only a stipulation, so my reasoning is weak, but I do so because my wider account focuses on comparing and matching claims by closeness, and so I prefer the reasoning behind *Match by Closeness*. It seems to me a slightly more respectful way to pair claims and slightly more parsimonious on the whole.

For those who are interested, the following is a version of my complete account which uses the *Match in Order of Strength* principle. If it later turns out the *Match in Order of Strength* is a better principle for comparing T-claims and F-claims, then we have a version ready to use.

The Complete Account (Match in Order of Strength Version)

- I. Identify the strongest remaining S-claim *with* competing S-claims [zero] levels lower. Match these S-claims and eliminate them from consideration
- II. Repeat step I until there are no S-claims with competing S-claims [zero] levels lower.
- III. Repeat steps I - II, with competing S-claims one level lower, then two levels lower, and so on until no S-claims with relevant competing S-claims remain.
- IV. Consider the remaining S-claims. Dismiss all S-claims that are not relevant to the strongest F-claims. I.e., dismiss all S-claims that do not have anchoring relevance.
- V. Dismiss from consideration any T-claims that are competingly irrelevant to any other competing T-claims.
- VI. Identify the strongest remaining T-claim. Match this T-claim with the strongest relevant F-claims and eliminate them from consideration.
- VII. Repeat step VI until:
 - a) there is a T-claim that cannot be outweighed by F-claims, in which case you should decide in favour of that group; or
 - b) there are no T-claims left, in which case use a fair procedure.

Appendix 5 – Competing Tiebreakers

In Chapter 6, I introduce the following step to my account:

- V. Dismiss from consideration any T-claims that are competing irrelevant to any other competing T-claims.

This step is introduced because without it we have cases with T-claims in both groups and it would be unclear how we ought to match those T-claims. This step means that we do not have competing tiebreakers in the two groups, and so that we only compare one group’s T-claims against the other group’s F-claims. To see the problem in detail, consider the following case:

Case 18

Level	Group A	Group B
1 - Death	± <1>	± <1>
2 - Paraplegia		1
3 - Lost Limb		
4 - Lost Hand		
5 - Lost Finger	1000	

First, let us assume that the lost finger claimants’ S-claims are not competing relevant to the paraplegia claimant’s S-claim in this case; the lost finger claimants’ S-claims could not aggregate to match or outweigh the paraplegia claimant’s S-claim. Nevertheless, let us also assume that both sets of claimants’ T-claims have anchoring relevance, such that they could tiebreak between the two deaths.

In this case, this means that we have competing tiebreakers. Now, the basic version of *Aggregate Relevant Tiebreakers* developed at the beginning of Chapter 5 simply tells us to weigh all the relevant T-claims together to see if they outweigh the F-claims. But this, of course will not work in heterogenous cases like this one. It would be strange to aggregate the paraplegia claimant’s T-claim and the lost finger claimants’ T-claims against the F-claims, given that they compete with each other. Instead, of course, we should compare only the T-claims in each group against the F-claims of the opposing group. Otherwise, T-claims in one group might be decisive in saving the opposing group!

Appendix 5 – Competing Tiebreakers

To do so, firstly, we need to recognise that without the death claimants and their F-claims we would save Group B outright because the lost finger claimants' S-claims are not relevant to the paraplegia claimant's S-claim. Thus, whilst the lost finger claimants' T-claims might be able to tiebreak between the death claimants, they cannot be the reason that the paraplegia claimant is not saved. For this reason, we dismiss the lost finger claimants' T-claims from consideration. So, in order to decide whether we should save Group B over running a lottery we should compare only the paraplegia claimant's T-claim against the death claimant's F-claim in Group A.

Appendix 6 – Local vs Global Relevance for F-claims

In Chapter 6, when consider Step VI of my sequence, I briefly note that I develop my view such that F-claims only need local relevance to competing T-claims to be able to outweigh the T-claims. I think it is worth addressing in a little more detail why we ought to use local relevance instead of global relevance when deciding whether an F-claim is relevant.

In simple, homogenous cases, where we have a single set of F-claims competing against one other set of T-claims, the problem of how we decide whether an F-claim is relevant does not arise.

However, there are many heterogenous cases in which there are a variety of different claims of differing strength. For instance, we might have the F-claims of dying people and people losing hands, compared to the T-claims of people facing paraplegia and those losing fingers. In such cases, it is not as easy to decide which F-claims and T-claims are relevant and we must decide whether to use global or local relevance to determine whether an F-claim is relevant or not:

Global Relevance: A claim, X, is Globally Relevant if X is strong enough compared to the strongest claim with which it competes.

Local Relevance: A claim, X, is Locally Relevant to another claim, Y, if X is strong enough compared to Y.

To see why this decision matters, let us reconsider the Case 17, this time split into two stages:

Case 17 - Stage 1

Level	Group A	Group B
1 - Death	10 <10>	10 <10>
2 - Paraplegia		
3 - Lost Limb		1
4 - Lost Hand		
5 - Lost Finger		1000

In this case let us assume that the single lost limb claimant's T-claim cannot by itself tiebreak between the death claims - the death claimant's F-claims outweigh the lost limb claimant's T-claim.

Appendix 6 – Local vs Global Relevance for F-claims

However, let us also suppose that the lost finger claimants' T-claims are relevant to the death claimants' F-claims. Thus, the lost limb and lost finger claimant's T-claims can together outweigh the F-claims and so they can tiebreak in favour of Group B.

Now consider Stage 2:

Case 17 - Stage 2

Level	Group A	Group B
1 - Death	10 <10>	10 <10>
2 - Paraplegia		
3 - Lost Limb		1
4 - Lost Hand	(1000) <1000>	(1000) <1000>
5 - Lost Finger		1000

At Stage 2, 1000 lost hand claims are added to both groups. These lost hand claimant's S-claims are matched with each other and taken out of consideration. However, they do, of course, have F-claims that might affect whether we should save Group B outright or hold a lottery. But how do we decide whether the lost hand claimants' F-claims are relevant to our decision? Do we compare the lost hand claimants' F-claims with the death claims, the lost limb claim, or the lost finger claims?

The first option is the Global Relevance option. We could compare the lost hand claimants' F-claims to the overall strongest claims with which they compete - i.e., the death claims. This will make the new F-claims irrelevant as death claims are so much stronger than lost hand claims. However, the number of death claimants is balanced, and so they only have F-claims. Of course, F-claims should not be made irrelevant by other F-claims, as they are not competing with each other. F-claims are all claims towards the same end (i.e., a lottery) so no-one with an F-claim would have any reason to object to the inclusion of other F-claims.

There is a second more moderate Global Relevance style option, however. We could compare the new F-claims to the strongest T-claim, in this case, the lost limb claim. The lost hand claimants' F-claims need not to be as strong to be relevant on this option as they do with the first global relevance option. This option also makes more sense than the first option - the lost limb claimant's

Appendices

T-claim is in direct competition with the F-claims. Thus, we might think that F-claims ought to be relevant to the lost limb claimant's T-claim if they are to weigh against saving Group B outright. Now, suppose we think that the lost hand claimants' F-claims are too weak to compete with the lost limb claimant's T-claim and should thus be considered irrelevant (as per *Stronger Always Tiebreaks* considered in the next appendix). This would mean that we are left with exactly the same consideration at Stage 2 as we had at Stage 1, and so we save Group B outright because of the lost limb and lost finger claimant's T-claims.

But this reasoning is not right. In this case, the lost limb claimant's T-claim is outweighed by the death claimants' F-claims. So, the lost limb claimant's T-claim is not decisive in favour of Group B. At Stage 1 it was the lost finger claimant's T-claims that eventually outweighed the death claimant's F-claims. Given that the lost limb claimant's T-claim is not decisive then it seems wrong to say that the lost hand claimant's F-claims need to be relevant to the lost limb claimant's T-claim to be taken into consideration.

Consider the justification we would need to give the lost hand claimants not to take their hands into consideration, and their objections. If we took the global relevance approach, we would have to tell the lost hand claimant's that we would not be taking into consideration their claims to an equal opportunity to be saved because we need to prioritise the lost limb claimant. However, the lost hand claimants would correctly object to this justification by pointing out that we are not prioritising the lost limb claimant, because the lost limb claimant would not otherwise be decisive between the two groups.

Nor does the lost limb claimant have a strong objection to the inclusion of these lost hand claimant's F-claims. The lost limb claimant is given a respectful justification as to why their claim is insufficient to tiebreak in favour of Group B, i.e., that there are death claimants who deserve an equal opportunity to be saved. Thus, the lost limb claimant cannot respectfully demand that they are given

priority over the lost hand claims when the lost limb claimant's claim has already been fully taken into consideration and outweighed.

This brings us on to our third option, the local relevance option. This option says we ought to compare F-claims only with the T-claims they need to outweigh. The local relevance option makes it easier for F-claims to reach relevance here than in either other option. F-claims only need to be relevant to the very claims they will outweigh. This seems like the right option, and it also follows more closely the reasoning we set out in Chapter 4.

In that chapter I argued that for a claim to tiebreak between other claims required that claim to have Global Relevance. I also argued that for a claim to be able to match or balance with another claim it only needed Local Relevance to that claim. I formulated this into *Anchoring* and *Competing Relevance*:

Anchoring Relevance: the relevance a claim needs to be able to become the new anchoring claim. A claim, X, has anchoring relevance iff X is sufficiently close in strength to the strongest claim (matched or unmatched) with which X is in competition, such that X can become the new anchoring claim when all stronger claims are matched.

Competing Relevance: the relevance a claim needs to be able to match or outweigh a competing claim. A claim, X, has competing relevance with another claim, Y, iff X is sufficiently close in strength to Y such that X can match with Y.

Here is a summary of my reasoning from Chapter 4. A claim that tiebreaks is not just competing with one particular claim, but it is competing with all the other claims. A claim that can be decisive in favour of one group is not just tiebreaking between the weakest claims but also tiebreaking between the strongest claims. Thus, a claim that is to be decisive must be relevant to all competing claims. The decisive claim must be able to form part of a respectful justification given to each of the claimants who is not saved.

Appendices

Conversely, when a claim is balancing another claim, it only needs to be relevant to that particular claim. A claim that balances with another claim need not be relevant to all competing claims because it is not denying those other claimants an opportunity to be saved. A claim that matches or balances with another claim is preventing that other claim from being decisive. It is thus only to that claim that it needs to be relevant.

With this in view, we can see that we ought to treat F-claims as only needing to be locally or competingly relevant. F-claims are by their nature unable to be decisive between groups. F-claims are only able to match and balance out other T-claims and prevent them from being decisive. Thus, they only need to be relevant to the T-claim which they are counterbalancing. In Case 17 this means that the lost hand claimants' F-claims only need to be relevant to the lost finger claimants' T-claims, and thus that we ought to hold a lottery at Stage 2.

Furthermore, no claimant could reasonably object to this approach of including and comparing claims. Each claimant is given a respectful justification as to why their claim is relevant or irrelevant and how their claims is outweighed by other relevant claims.

Appendix 7 – When are F-claims irrelevant?

When comparing F-claims and T-claims to decide whether we ought to run a lottery or save a group outright, we need to know when the F-claims and T-claims are irrelevant and should be dismissed from consideration. The simple formulation of my view does not give any guidance on when an F-claim is irrelevant to a competing T-claim. However, it is vital that F-claims are considered irrelevant in some cases. Otherwise, many weak F-claims could outweigh much stronger T-claims as per the *Strong Tiebreakers* case:

Strong Tiebreakers: You can save Group A or Group B. In Group A you have N number of sore throat claims and in Group B you have N number of sore throat claims and a death claim.

Thus, we face the question of how to decide whether an F-claim is relevant to the particular T-claims with which it is in competition - i.e., under what conditions are F-claims irrelevant to particular T-claims?

In Chapter 6, when developing Step VIII of my view, I state that F-claims are only relevant to T-claims with the same strength, or weaker, underlying stake. I.e., a paraplegia claimant's F-claim is relevant to other paraplegia claimant's T-claims, but not to a death claimant's T-claim. I call this principle *Stronger Always Tiebreaks*:

Stronger Always Tiebreaks: A T-claim always tiebreaks between F-claims of those with less at stake. I.e., F-claims are irrelevant to the T-claims of claimants with more at stake.

I come to this position through a process of reflective equilibrium, based on some of my intuitions, some considered convictions and how it fits with wider limited aggregation justifications.

Firstly, let us consider our intuitive reasons for *Stronger Always Tiebreaks*. My strong intuition is that F-claims are relevant to T-claims of the same strength or weaker. *Kamm's Large Scale Case* highlights this intuition:

Appendices

Kamm's Large Scale Case: You can save Group A or Group B. In Group A you have 1000 people who will die. In Group B you have 1001 people who will die.

In this case, intuition suggest that we ought to run a lottery between the two groups. Now consider another large-scale case:

Large Scale Case 2: You can save Group A or Group B. In Group A you have 10,000 people who will be paralysed. In Group B you have 10,000 people who will be paralysed and another who will die.

My intuition in this case is that we should save Group B. It strikes me that the death claimant has the most to lose and that the F-claims of those with less at stake are not relevant to her T-claim.

Admittedly, however, my intuition is significantly weaker in this case, and so I do not feel sufficiently confident in setting relevance at this level on this basis.

However, I do think that there are also good theoretical reasons to set relevance at this level. Firstly, limited aggregation views try to balance the concerns of the aggregators and the anti-aggregators. I think to say that we ought to run a lottery in *Large Scale Case 2* moves us too far away from the anti-aggregative position that we ought to save the strongest claimant. To say that a stronger claim could not tiebreak between weaker claims seems strange to say the least. It does not seem to respect the death claimant's priority over the other claimants.

Remember that limited aggregation views are based around giving due consideration to each individual's partial concern and respecting the personal perspective of each of the claimants.⁵⁹ The easiest way to capture this is through Bastian Steuwer's *Respectful Failure to Save Principle*:

⁵⁹ As I have made clear in previous chapters (see Chapter 3 Section 3) I do not want to commit myself to any particular limited aggregation view but shall develop my views in such a way that they appeal in a more general way. This is an advantage throughout most of the thesis, however it is a bit of a disadvantage here as it is harder to show exactly why limited aggregation views should prioritise the partial concern of the death claimant over the concerns of the paraplegia claimants. Nevertheless, I hope what I say here gives a sufficient overview as to why we ought to save Group B in *Large Scale Case 2*.

Appendix 7 – When are F-claims irrelevant?

“Respectful Failure to Save Principle. Every person whom we fail to save is entitled to a respectful justification for our failure to save. It is disrespectful and impermissible to fail to save a person with a strong claim for the sake of persons whose claims are irrelevant to this strong claim.” (Steuwer, 2021a, p. 24)

In *Large Scale Case 2* it seems like we can give the paraplegia claimants an acceptable justification for why we do not give them an equal opportunity to be saved: “I’m sorry that we cannot run a lottery to give you an equal opportunity to be saved. There is somebody’s life on the line, and we must save them.”. Such a justification would respect the paraplegia claimant, and he could not reasonably give any kind of objection.

On the other hand, if we were to give the same justification to a death claimant in *Kamm’s Large Scale Case*, that death claimant might reasonably respond: “What about my life though? Is my life not as valuable as the other person’s? Do I not deserve at least an opportunity to be saved?”. Now, we might still be able to give this claimant a respectful justification for not giving them an opportunity to be saved, if the T-claims are strong enough, but at least by considering the death claimant’s F-claim we are giving the claimant due concern.

Similarly, we might think it disrespectful to the death claimant to consider the paraplegia claimants’ F-claims. The death claimant might reasonably argue that his T-claim is stronger than the paraplegia claimants’ F-claims, and that given that we can save an equal number of paraplegia claimants whomever we save, to consider their F-claims when his S-claim is on the line is thus disrespectful to him and what he has to lose.

Secondly, we should consider what pro tanto reasons we have. In Chapters 2 and 3 we discussed how relevance views treat the strongest claims as having a pro tanto claim to be saved. This claim must be outweighed if we are not to save the strongest claimant. The strongest claimant needs a special justification if we are not to satisfy her claim.

Now, other people's S-claims can form part of that special justification so long as they are strong enough. But it seems that other people's F-claims could not form part of this justification. It is one thing weighing the S-claims of weaker claimants against her claim, but it is an entirely different thing to weigh those claimant's F-claims against her claim. I think the concept of dominance, and in particular strict dominance, is very helpful here:

Strict Dominance: A claim, or set of claims, strictly dominates another claim, or set of claims, if it is stronger in all respects.

In *Large Scale Case 2* the death claim strictly dominates each paraplegia claim. The death claim is stronger both in terms of what is at stake for them, and in terms of type of claim (i.e., a T-claim is a stronger type of claim than an F-claim).⁶⁰ It strikes me that this is a good place to draw the line for relevance in these cases - a strictly dominated claim should not be relevant to the claim which dominates it.

Thus, weaker types of claims (F-claims) can sometimes aggregate against stronger types of claims (S-claims) if what is at stake for those with the weaker type of claim is greater or equal to what is at stake for those with the stronger type of claim. Similarly, claimants with less at stake can sometimes aggregate their claims against those with more at stake, on the condition that the type of claim they have is stronger or equal to the type of claim that those with more at stake have. However, a weaker type of claim of those with less at stake cannot aggregate against a stronger type of claim of someone with more at stake.

⁶⁰ One might think that the death claim does not strictly dominate the paraplegia claims, because the paraplegia claims are greater in number. However, as we have discussed in previous chapters relevance does not take into consideration numbers in this way. Relevance, as an anti-aggregative notion, respects the separateness of persons and uses the personal perspective or partial concern of individuals to determine relevance. Thus, we must consider each paraplegia claimant's relevance in comparison with the death claim individually - and in *Large Scale Case 2* the paraplegia claims are strictly dominated.

Appendix 8 – Fair procedure

My thesis, and the account developed in Chapter 6, does not address what exactly a fair procedure requires of us. Do we simply flip a coin between the two options, or do we run a weighted lottery, and if so, which claims determine the weight of the lottery? For reasons of space, I have not addressed this in the main thesis, but I want to briefly address this issue and demonstrate how my view can accommodate either approach.

The first, and simplest, option is to flip a coin between the two groups, whenever we ought to do what is fair. Flipping a coin is certainly a better option than choosing one group outright in some of these cases. In fact, in some cases, such as *Kamm's Sore Throat case*, flipping a coin seems like the right answer.

However, there are other cases where this does not seem to be the correct approach. Consider a case with 10 dying people in Group A and 12 people facing dying in Group B. Or suppose we have 10 dying people in Group A and 30 people facing paraplegia in Group C. Whilst it might be right to use some fair procedure to determine whom to save, we might think that flipping a coin does not capture the proportional strength of claims in these cases. In these cases it seems like a weighted lottery does a better job of capturing the right thing to do.

Using a weighted lottery might better fit with the deeper justification for this view, explored in Chapter 6, too. In that chapter we looked at how *Aggregate Relevant Tiebreakers* combines Kamm's priority approach with Broome's proportionality approach to fairness. Now according to Broome "*claims should be satisfied in proportion to their strength.*" (Broome, 1990, p. 95) and when they cannot all be satisfied, they ought to be given an opportunity to be saved proportional to their strength. If this is the reasoning behind *Aggregate Relevant Tiebreakers*, then to flip a coin in these cases might not always be the fairest thing to do.

Now, I do not want to debate the merits of a weighted lottery approach. I do not think there is much benefit in covering this well-trodden ground in the remainder of the thesis. But I do want to quickly explain how a weighted lottery style approach can be combined with my overall view.

Suppose then that the fair procedure requires us to run a weighted lottery of some sort. However, we still have the question about exactly which claims should be taken into consideration when determining the weight of the lottery. The following are several potential principles for deciding which claims are relevant to determining the weight of a lottery, ordered from least to most restrictive:

All Claims: The weighted lottery should be determined by the aggregated weight of all claims.

Anchoring Relevant Claims: The weighted lottery should be determined by the aggregated weight of only those claims with anchoring relevance to the strongest claim.

Competing Relevant Claims: The weighted lottery should be determined by the aggregated weight of only those claims with competing relevance to the strongest tiebreaking claim.

Anchoring & Competing Relevant Claims: The weighted lottery should be determined by the aggregated weight of only those claims with anchoring relevance to the strongest claim and competing relevance to the strongest tiebreaking claim.

Strongest Claims: The weighted lottery should be determined by comparing the proportional strength of the strongest claim in each group, ignoring the quantity of claims.

Firstly, I think it is worth looking at the most restrictive option - *Strongest Claims*. This seems to be the account that Broome takes. His argument for weighted lotteries relies entirely on the difference in strength between claims and not their numbers, this is why he thinks that a fair procedure in the case where we can save 1 person from dying or 5 people from dying is to give both groups 50%

chance, and that the 5 cannot pool their chances to have a stronger chance of winning (Broome, 1998, pp. 4-5). This option is thus different to the other options by ignoring the numbers of claims.

However, I think this is a mistake of Broome's account and I think any weighted lottery ought to be sensitive to numbers. Firstly, on a purely intuitive level, it seems fairer to me that a group of 12 dying people have a greater chance of being saved than the group of 10 dying people. Of course, Broome would not accept this on my intuition, but I think it is a good starting point.

More pertinently, *Strongest Claims* has some strange and unpalatable consequences. First, suppose we could save (A) 100 people from dying or (B) 1 person from dying and 1 million from sore throats. On Broome's view, without the sore throats the utility of saving (A) outweighs the requirements of fairness. But with the 1 million sore throats the utility of saving (A) is not enough to outweigh the requirements of fairness. But because the lottery is determined only by the strongest claims, we ought to give both groups an equal 50% chance of the good. Thus, the sore throats, despite only deserving a very small proportional chance of the good, end up supporting a much weightier lottery.⁶¹

Now on my account the sore throats would be dismissed from consideration early, so this extreme consequence of *Strongest Claims* can be avoided on my view. But equivalent cases with smaller disparities in claims will lead to similarly bizarre consequences. For instance, the addition of some broken finger claims might tip the scales in favour of a lottery by counterbalancing with some other weaker claims. However, the weight of the lottery would then be determined by the strongest claims in the competition. Thus, the broken finger claims despite only deserving a small proportional chance of the good, end up supporting a much weightier lottery in their favour.

Similarly, we could be deciding between two groups A and B. Let us suppose that Group A is overall the stronger group, but not quite strong enough to mean we ought to save Group A outright (the

⁶¹ There are other problematic consequences of *Strongest Claims* for Broome, but this should be sufficient for our purposes.

exact claims do not matter). Thus, we ought to run a lottery. However, Group B has the strongest claim, and thus the lottery according to *Strongest Claims* ought to favour (and potentially significantly favour) Group B, despite the fact that Group A is the overall stronger group. I think these are sufficient reasons to dismiss *Strongest Claims* and to prefer an account that does consider the number of claims.

Thus, let us turn to the options that include the numbers in the weighting and let us start at the other end of the spectrum and look at the least restrictive option – *All Claims*. This option is very simple - it tells us to weigh all the claims in both groups and determine the weight of the lottery by this weight. However, *All Claims* allows very weak claims to sometimes overwhelm stronger claims. To see why, consider the following case:

Case 19

Level	Group A	Group B
1 - Death	10 <10>	10 <10>
2 - Paraplegia	1	
...		
100 – Sore Throat		N

In this case, let us suppose that the paraplegia claim is not strong enough to tiebreak between the two groups of death claims. *Aggregate Relevant Tiebreakers* thus tells us to run a weighted lottery. But rather than having a weighted lottery slightly in favour of Group A, *All Claims* tells us to incorporate the very large number of sore throats in Group B into our decision. Thus, our weighted lottery might end up massively in favour of Group B, because the sore throats (in large enough numbers) overwhelm all the other considerations. This, of course, is the exact sort of consequence that limited aggregation views are designed to avoid. Thus, we can dismiss *All Claims* from consideration.

Next, let us look at *Anchoring Relevant Claims*. This option tells us to only consider claims which could theoretically tiebreak between the strongest claims if there were enough of them. This option prevents cases like Case 19 above and makes a lot more theoretical sense. It seems reasonable to

Appendix 8 – Fair procedure

limit the claims which can affect the weight of the lottery to those claims which are relevant to the strongest claims. Otherwise, we are letting claims irrelevant to the strongest claims effect the chances of saving the strongest claims.

Nevertheless, this option has a few problems of its own. Firstly, if anchoring relevance is determined by competition (i.e., a claim is anchoringly relevant if it is strong enough to the strongest *competing* fairness claim), then we will get unusual cases like the following:

Case 20

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia		2
3 - Lost Limb		
4 - Lost Hand		
5 - Lost Finger		
6 - Broken Arm	1000	1000

In this case, the broken arms in Group A are relevant but the broken arms in Group B are irrelevant.

Thus, *Anchoring Relevant Claims* tells us to take the Group A broken arms into consideration, but not the Group B broken arms when determining the weight of the lottery. This means that the weighted lottery will hugely favour Group A, in a way that seems unfair. If the broken arms in Group A can affect the chances of which group will be saved, then so should the broken arms in Group B.

Otherwise we violate *Equal Consideration for Equal Claims*:

Equal Consideration for Equal Claims: all claims of equal strength ought to be given equal weight in determining which group to save. (Tomlin, *On Limited Aggregation.*, 2017, p. 243)

Perhaps we ought to determine anchoring relevance by strength instead. This would mean determining whether a claim has anchoring relevance by comparing it to the strongest claim in the competition as a whole. For instance, in this case by comparing the broken arm claims with the death claim. This would solve the issue in Case 20, as it would mean that none of the broken arm claims count.

Appendices

However, if we take this approach we will still end up with another problematic case. Consider the following:

Case 21

Level	Group A	Group B
1 - Death	10 <10>	10 <10>
2 - Paraplegia		1
3 - Lost Limb		
4 - Lost Hand		
5 - Lost Finger	1000	

In this case, the thousand lost finger claims are irrelevant compared to the paraplegia claim, and so are taken out of consideration. However, suppose that the paraplegia claim cannot tiebreak between the death claims and so we must use a fair procedure. *Anchoring Relevant Claims* tells us to include all the lost finger claims in determining the weights of this lottery and so the lottery will heavily favour Group A. But this does not seem right given that the lost finger claims are irrelevant to the paraplegia claim.

Next, let us consider *Competing Relevant Claims*. *Competing Relevant Claims* tells us to weigh only those claims that are relevant to the strongest T-claim. This makes theoretical sense in a similar way to *Anchoring Relevant Claims*. It seems reasonable to limit the claims which can affect the weight of the lottery to those claims which are relevant to the strongest tiebreaking claim. Otherwise, we are letting claims irrelevant to the strongest T-claim effect the chances of the T-claim being saved. Of course, because the T-claim is an S-claim, relevance should be set by competing relevance here instead of anchoring relevance.

Now, *Competing Relevant Claims* does not have the problem of the previous views. In Case 19, the sore throats are irrelevant to the paraplegia claim and so do not factor into the weighting of the lottery. Similarly, in Case 21, the lost finger claims are irrelevant to the paraplegia claimant's T-claim, and so the lost finger claims do not factor into our decision. In Case 20, there is no tiebreaking claim

Appendix 8 – Fair procedure

and so the broken arm claims in both groups remain relevant, and we are not treating the two groups differently.

However, *Competing Relevant Claims* also has problems. Consider the following case, very similar to Case 19:

Case 22

Level	Group A	Group B
1 - Death	10 <10>	10 <10>
2 - Paraplegia		
...		
100 – Sore Throat		N

In this case, the strongest T-claims are the sore throat claims. Trivially, sore throats are relevant to themselves. This means that *Competing Relevant Claims* does not have a way to tell us that the sore throat claims should not be taken into consideration when determining the weight of the lottery. Thus, *Competing Relevant Claims* will include the sore throats when determining the weights of the lottery, and thus allow the sore throat claims to overwhelm the death claims, such that we ought to favour Group B significantly in the lottery. This, of course, is also wrong. The whole point of relevance views is to prevent such consequences.

Lastly, let us look at *Anchoring & Competing Relevant Claims*. This approach combines the two previous views. This means that we only consider claims that are anchoringly relevant to the strongest claims and competingly relevant to the strongest tiebreaking claim. This view thus combines the theoretical bases of the two previous views. To be able to affect the chances of the strongest claimants being saved, a claim must be anchoringly relevant to that claim, and to be able to affect the chances of the tiebreaking claims being saved, a claim must be competingly relevant to that claim.

Anchoring & Competing Relevant Claims can also, therefore, capture the above cases. In Case 19 and Case 22 *Anchoring & Competing Relevant Claims* tells us that the sore throat claims are not

Appendices

anchoringly relevant to the strongest claims and so do not factor into the weighting of a lottery. In Case 21, the lost finger claims are not relevant to the paraplegia claim and so, again, do not factor into the weighting of a lottery.

However, we might still have issues with Case 20. Claims of the same strength on both sides will need to be included in the competition if we are to avoid violating *Equal Consideration for Equal Claims* in such cases. The solution to this problem is fairly simple - we can easily stipulate that claims of the same strength must be treated equally. To do so, we simply state that all claims must be anchoringly relevant to the strongest claim and competingly relevant to the strongest T-claim, regardless of whether they compete with that claim or not.

With this in view, we can see how my complex relevance view could be combined with a lottery-based approach to fairness and how we would determine the weights of such a lottery. You might disagree with my conclusion about the best approach to weighted lotteries or even whether lotteries ought to be weighted. But what I hope I have shown here, is that my wider view is at least compatible with weighted lotteries.

Appendix 9 – Tiebreaking and the Principle of Net Addition

In Chapter 1 I demonstrate that any relevance view will violate *Weak Additivity*:

Weak Additivity: If A is preferable to B, and C is preferable to D, then A with C is preferable to B with D.

I argued that violating this principle should not be considered a problem with the view but a necessary component of how we ought to think about these cases. However, I nevertheless argued that not all violations of *Weak Additivity* are made equal - and that some violations are not acceptable. One of these was the *Principle of Net 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. (Van Gils & Tomlin, 2019, p. 253)

However, *Aggregate Relevant Tiebreakers*, and any other view that takes fairness seriously, will seem to violate this principle. To see why consider this case:

Case 23

Level	Group A	Group B
1 - Death	1 (+10,000)	2 (+10,001)

In this case at Stage 1 we obviously save Group B. But at Stage 2 with the addition of a huge number of death claims to each group (but ever so slightly more to Group B) we switch away from saving Group B and run a lottery instead. In this way it seems that Group B has become less choice-worthy despite having more claims of the same strength added. This might seem like a major problem for the view then – it violates a fairly fundamental principle of aggregation. Either that, or we have to abandon the *Principle of Net Addition* which has been a part of the motivation for this view.

I think the first thing to say about this problem is that this is not just a weird consequence of the view, or a hidden problem for views that take fairness seriously. Cases like Case 23, where we have

10,001 people in one group and 10,003 in the other, are fundamental to limited aggregation views (and even fully aggregative views, such as Broome's and Hirose's) which try to accommodate considerations of fairness. The intuition, that even aggregators like Broome and Hirose share, is that we ought to give both groups in cases like these fair chances. Given the strength of the intuition in this case, and the centrality to such views, it strikes me that if there is a clash between the *Principle of Net Addition* and our intuition in these cases, then we should abandon the *Principle of Net Addition*, or at the very least put limits on the principle.

Doing so, however, will seriously undermine much of what I have said throughout the thesis and undermine one of the main ways with which we can analyse different limited aggregation cases. Plus, the initial plausibility of the *Principle of Net Addition* makes biting the bullet on the principle quite difficult and will make limited aggregation views significantly less appealing. The best potential outcome is one in which we can maintain the principle and capture our intuition in these cases too. Thus, I hope to quickly show that we do not actually violate the *Principle of Net Addition* here, even though it seems like we do.

With that in mind, I think it is helpful to recognise that in these cases there are not two choices for what we ought to do, but three. We can save Group A, Group B, or use a fair procedure. Once we recognise this, we can see that Group B has not become less choice-worthy compared to Group A with the addition of the large number of extra claims to both groups. Instead, Group B has become less choice-worthy compared to the third option, that of using a fair procedure. Once we recognize this the problem disappears.

Whilst, throughout the thesis, I have put the F-claims in the same box as the claimants' S-claims in our tables. This is not the only way I might have drawn the tables. Now, I think the way I have drawn the tables is the most helpful and intuitive way to understand how these claims work (largely because we only count one side's F-claims). However, we could alternatively have expressed cases like Case 23 like this:

Appendix 9 – Tiebreaking and the Principle of Net Addition

Case 23

Level	Group A	Group B	Fair Procedure
1 - Death	1 (+10,000)	2 (+10,001)	1 (+10,000)

Once we express our options like this, we can see that the choice-worthiness of Group B at Stage 2 has increased relative to Group A by a single death claim, to two death claims. So, the net choice-worthiness of Group B over Group A is now two death claims. However, at the same time, the third option of running a lottery has increased its choice-worthiness by significantly more. Running a lottery has increased its choice-worthiness from a single F-claim to 10,001 F-claims. At Stage 1, we had to compare one death claimant's T-claim to one death claimant's F-claim. At Stage 2, we had to compare two death claimant's T-claims to 10,001 death claimants' F-claims. Hopefully, this should make it obvious why my account does not violate the *Principle of Net Addition*.

Appendix 10 – Aggregate Relevant Tiebreakers and Expected Value

In Chapter 6 I develop Aggregate Relevant Tiebreakers by comparing the strength of a T-claim against the strength of an F-claim. I do not adjust the strength of these claims according to expected value. To make this clear, consider the following case:

Case 24

Level	Group A	Group B
1 - Death	10 <10>	10 <10>
2 - Paraplegia		1

In this case, my sequence compares the paraplegia claimant's T-claim directly with the death claimants' F-claims in Group A. But we could compare the claims by adjusting their strength by the expected value of each option.

Suppose we only have two options: flip a coin or save Group B outright. If we flip a coin, the paraplegia claimant still has a 50% chance of being saved. Whereas if we save Group B outright, the paraplegia claimant is guaranteed to be saved. Thus, we might think that we should adjust the strength of the T-claim down to 50%: the paraplegia claimant's T-claim is less strong given that she has a good chance of being saved anyway. Similarly, we might think that we ought to adjust the strength of the death claimants' F-claims according to expected value. If we choose to save Group B outright, the death claimants only lose a 50% chance of being saved.

Importantly, the expected value of the options might change depending on what a fair procedure requires. For instance, in Case 24 suppose the fair procedure is not to flip a coin but to run a weighted lottery, such that we ought to give Group A a 45% chance of being saved and Group B a 55% chance. If this is the case, then the additional value of saving Group B outright is reduced – Group B already has a 55% chance of being saved so the additional utility of saving Group B over running the lottery is only 45% extra chance of saving the paraplegia claim. Thus, if we run the lottery the paraplegia claimant only loses out on a 45% extra chance of being saved.

Similarly, we would need to adjust the strength of the death claimants' F-claims. If we choose to save Group B outright, the death claimants in Group A only lose out on a 45% chance of being saved. So, the strength of the F-claims is also reduced. Importantly, this means that including expected value does not alter how we weigh T-claims against F-claims. If we adjust for expected value, then we reduce the strength of both the T-claims and the F-claims by the same proportion and so the reductions cancel each other out. Thus, for simplicity we might as well not adjust for expected value.

Nevertheless, if one wants to amend the account to include adjusting for expected value one need only add a step to *Aggregate Relevant Tiebreakers* and the Complete Account which adjusts the value of the T-claims and F-claims respectively. This step would simply require determining how the lottery ought to be weighted and then adjusting the strength of the respective claims. If we use the principle for determining which claims affect the lottery outlined in Appendix 8, then the following is such a step that could be added after Step 3 of *Aggregate Relevant Tiebreakers* or after Step V of the completed sequence:

Expected Value Step: Identify all the S-claims with anchoring relevance to the strongest S-claim and competing relevance to the strongest T-claim. Aggregate their weights to determine the weight of the lottery. Adjust the strength of the F-claims and T-claims accordingly.

Other versions of this step could easily be generated for different ways of calculating a weighted lottery. The flexibility of the account extends even further, such that if one were to find a reason to treat F-claims and T-claims disjunctively (such that only one type of claim is amended by expected value but the other is not – or that they are amended by different proportions), then a version of this step could also be generated to adjust for that disjunct. I cannot think of any reason why we would treat the claims disjunctively, but it is nevertheless an advantage of the account that it can cope with such approaches.

Appendix 11 – Partial Leftovers

One assumption I make throughout the thesis is that claims match exactly. This is an assumption I make primarily for simplicity: dealing with proportions of partially counterbalanced claims is adding an unnecessary level of detail and will easily confuse an already complicated account. However, whilst this assumption is made for simplicity it does skip over an important ambiguity. Namely, how do we determine if a claim is relevant to a partially matched claim? For instance, we might think that a lost hand claim is irrelevant to a death claim, but is a lost hand claim irrelevant to half a death claim when the death-claim is partially outweighed by a paraplegia claim? What about to a tenth of a death claim?

Van Gils and Tomlin, who identify this issue, offer two potential approaches to such cases:

Full Claim Relevance: To be relevant, a claim must be relevant to the claim with which it competes.

Partial Claim Relevance: To be relevant, a claim must be relevant to the proportion of the claim with which it is in competition. (Van Gils & Tomlin, 2019, p. 249)

They both prefer *Partial Claim Relevance*. They argue that it would be unfair to the weaker claimant to have to compare their claim to the full-strength stronger claim when the stronger claim is already partially matched. The weaker claim only needs to compete with a proportion of the stronger claim, and thus should only need to be relevant to that proportion.

I think Van Gils and Tomlin are wrong here, however. Firstly, as they note, this account would allow very weak claims to be relevant and to be decisive in favour of one group. They note that this also gives us further reason to step away from the sore throat intuition in tiebreak cases. Yet, I think the sore throat intuition is much more fundamental to limited aggregation views than Van Gils and Tomlin acknowledge. Given that the hybrid account I develop in the thesis can capture the sore throat intuition (unlike Van Gils and Tomlin's view) I do not think we have good reason to abandon

it. If *Partial Claim Relevance* is inconsistent with the sore throat intuition, then *Partial Claims Relevance* is the one that ought to go.

Secondly, I find the reasoning Van Gils and Tomlin give in defence of *Partial Claims Relevance* problematic. If a weaker claim only needs to be relevant to a proportion of a stronger claim to be able to compete with that claim, then it seems that this has the potential to devolve into a fully aggregative view. If we say that each of a very large number of weak claims only competes with a small proportion of a stronger claim, then we can say that all those claims are relevant to the competition.

For instance, suppose we can save one person from dying or 100,000 from a sore throat. We might say that each of the sore throat claimants are only competing with $1/100,000^{\text{ths}}$ of a death claim, and that each sore throat claim is relevant to that very small proportion. Thus, all the sore throats are relevant and they on mass outweigh the death claim. This, of course, is the exact opposite conclusion the limited aggregation views are designed to avoid.

Now, one might argue that the sequential nature of Van Gils and Tomlin's view prevents this problem. We could say that it is not about what proportion of a stronger claim a weaker claim needs to compete with, but what proportion of a stronger claim actually remains unmatched when other relevant claims have been considered. However, I find this response somewhat artificial. If a claim needs only to be relevant to a proportion of a claim, I do not see why it matters whether that proportion is a part of a partially matched claim or not.

Furthermore, we have some odd consequences with *Partial Claims Relevance*. Consider this case:

Appendices

Case 25 - Stage 1

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia		
3 - Lost Limb		
4 - Lost Hand		1
5 - Lost Finger		1
6 - Broken Arm		1
...		...
1000 - Sore Throat		1

At Stage 1, none of the claims in Group B are relevant to the death claim in Group A. Thus, we save Group A outright.

But now consider Stage 2. At this stage, the infection that meant the claimant in Group B who would lose a hand has now spread further up his arm. He will now lose his whole limb:

Case 25 - Stage 2

Level	Group A	Group B
1 - Death	1	
2 - Paraplegia		
3 - Lost Limb		1
4 - Lost Hand		
5 - Lost Finger		1
6 - Broken Arm		1
...		...
1000 - Sore Throat		1

Of course, this now means his lost limb claim is relevant to the death claim. However, the slight strengthening of this claim, not only makes the lost limb claim relevant, but it also makes all the other Group B claims relevant to the death claim.

This is because the lost limb claim matches with the death claim leaving the death claimant with only part of their claim left. The lost finger claim is then relevant to the proportion of the death claim that remains after the lost limb claim is matched, and so can match with it. This in turn reduces the proportion of the death claim further such that the broken arm claim is relevant and can match. This

process carries on until the sore throat claim is relevant to the remaining proportion of the death claim, which it might even outweigh.

This strikes me as a profoundly odd conclusion. The slight strengthening of one claim should only make that claim relevant to the competing claim. That one patient's infection has spread up their arm ought not to mean that a sore throat is suddenly relevant to a death claim.

On a limited aggregation view whether a claim is relevant to the overall competition ought to depend on the other claims in competition, but whether that claim is relevant to a particular competing claim ought not to depend on the other claims in the competition. This is an important distinction to make.

It is this reasoning that also convinces me that we ought to use *Full Claims Relevance* when dealing with partial leftovers. A claim's relevance to a competing claim ought to rely entirely on what is at stake for those two claimants and not on what other claims are in the competition. A sore throat cannot be a reason not to save a life, whether it is in conjunction with many other sore throats or with an assortment of stronger claims.

Consider the justifications behind limited aggregation views. Again, I do not commit myself to a particular account, but all the accounts share a concern for the partial concern of each individual and what is at stake for them. It is the individual's stake which makes other claims irrelevant. The stake does not change when part of the claim is matched: an individual with a death claim either lives or dies.

Thus, a weaker claim might come into conflict with a partially matched claim, but that weaker claim does not come into conflict with a partially outweighed stake. What a person might lose is not less simply because part of their claim is matched. The individual's interest is in their stake not their claim.

Appendices

To put it in Steuwer's terms (Steuwer, 2021a, p. 24) a weak claim cannot form part of a justification for our failure to save someone with a much stronger claim. It would be disrespectful to a dying person to say that we cannot save them due to a sore throat claim. This is true whether that sore throat claim is combined with paraplegia claims or with other sore throat claims. Similarly, in Vooehoeve's terms (Voorhoeve, 2014) it would be impermissible for a sore throat claimant to prefer her claim be treated instead of someone who is dying and, again, this does not depend on what other claims are in the competition.

I think the above is sufficient to decide in favour of *Full Claims Relevance*. The rest of my view is made with this assumption in mind. Thus, in cases where a claim is partially matched, we carry on the sequence as usual – my account remains unchanged. Relevance stays the same, but the number of claims needed to match a partially matched claim is reduced. For instance, consider this case.

Case 26:

Level	Group A	Group B
1 - Death	1 [0.5]	
2 - Paraplegia		1
3 - Lost Limb		2

First, we match the paraplegia claim with the death claim such that half the death claim is matched. The remaining half a death claim is then matched by the two lost limb claims. We thus have a tied case.

Perhaps, the reader prefers *Partial Claims Relevance* and they have some strong reason to prefer that approach. If so, the account I develop in this thesis will need significant altering to capture this. I do not think it would be impossible to design such an account; however, this account would be significantly more complicated, and so I shall not design one. I leave it to the reader to generate a *Partial Claims Relevance* version of my account if they so desire.

Bibliography

- Broome, J. (1990). Fairness. *Proceedings of the Aristotelian Society*, 91, 87-101.
- Broome, J. (1991). *Weighing Goods: Equality, Uncertainty and Time*. Oxford: Blackwell.
- Broome, J. (1998). Kamm on Fairness. *Philosophy and Phenomenological Research*, 58, 955-961.
- Broome, J. (2002). All Goods are Relevant. In C. J. Murray, J. A. Salomon, C. D. Mathers, & A. D. Lopez, *Summary Measures of Population Health* (pp. 727-729). Geneva: World Health Organization.
- Broome, J. (2004). *Weighing Lives*. Oxford: Oxford University Press.
- Daniels, N. (2020, Summer). *Reflective Equilibrium*. (E. N. Zalta, Ed.) Retrieved from Stanford Encyclopedia of Philosophy:
<<https://plato.stanford.edu/archives/sum2020/entries/reflective-equilibrium/>>
- Eggel, M., & Martin, A. K. (2022). Limited Aggregation for Resolving Human-Wildlife. *Ethics, Policy & Environment*,.
- Ethox Centre. (2023, January 9). *Healthcare resource allocation in the aftermath of the Covid-19 pandemic*. Retrieved from Ethox Centre Website: <https://www.ethox.ox.ac.uk/Our-research/research-projects/healthcare-resource-allocation-aftermath-covid-19>
- Greaves, H., & MacAskill, W. (2021). The case for strong longtermism. *GPI Working Paper*.
- Halstead, J. (2016). The Numbers Always Count. *Ethics*, 789-802.
- Hirose, I. (2015). *Moral Aggregation*. Oxford: Oxford University Press.
- Hooker, B. (2005). Fairness. *Ethical Theory and Moral Practice*, 8, 329-352.
- Horton, J. (2018). Always Aggregate. *Philosophy & Public Affairs*, 46(2), 160 - 175.
- Horton, J. (2020). Aggregation, Risk and Reductio. *Ethics*, 514-529.
- Kagan, S. (2019). *How to Count Animals, More or Less*. Oxford: Oxford University Press.
- Kamm, F. (1998). *Morality, Mortality Volume I: Death and Whom to Save From It*. Oxford: Oxford University Press. doi:10.1093/0195119118.001.0001
- Kamm, F. (2002). Health and Equity. In C. J. Murray, J. A. Salomon, C. D. Mathers, & A. D. Lopez, *Summary Measures of Population Health* (pp. 685 - 706). Geneva: World Health Organization.
- Kamm, F. (2007). *Intricate Ethics: Rights, Responsibilities, and Permissible Harm*. Oxford: Oxford University Press. doi:10.1093/acprof:oso/9780195189698.001.0001
- Kamm, F. (2015). Bioethical Prescriptions. *Journal of Medical Ethics*, 41(6), 493-495.
- Kelleher, J. P. (2014). Relevance and Non-Consequentialist Aggregation. *Utilitas*, 26(4), 387-408.
- McMahan, J. (2009). Asymmetries in the Morality of Causing People to Exist. In M. A. Roberts, & D. T. Wasserman, *Harming Future Persons* (pp. 49-68). Dordrecht: Springer.

- Mogensen, A. (2019). Staking our future: deontic long-termism and the non-identity problem. *GPI Working Paper*.
- Murray, C. J., & Lopez, A. D. (1996). *The Global Burden of Disease*. World Health Organisation.
- Nagel, T. (1995). *Equality and Partiality*. Oxford: Oxford University Press.
- NHS England. (2017, November 30th). *Prescription curbs to free up hundreds of millions of pounds for frontline care*. Retrieved from NHS England: <https://www.england.nhs.uk/2017/11/prescription-curbs-to-free-up-hundreds-of-millions-of-pounds-for-frontline-care/>
- Norcross, A. (2002). Contractualism and Aggregation. *Social Theory and Practice*, 303-314.
- Perry, P. A., & Hotze, T. (2011). Oregon's Experiment with Prioritizing Public Health Care Services. *Virtual Mentor*, 241-247.
- Piller, C. (2017). Treating Broome Fairly. *Utilitas*, 214-238.
- Rachels, S. (1998). Counterexamples to the transitivity of better than. *Australasian Journal of Philosophy*, 76(1), 71-83.
- Ross, W. (1930). *The Right and the Good*. Oxford: Oxford University Press.
- Rüger, K. (2020). Aggregation with Constraints. *Utilitas*, 32, 454-471.
- Scanlon, T. (1998). *What We Owe to Each Other*. Cambridge, MA: Harvard University Press. doi:10.2307/j.ctv134vmrn
- Singer, P. (1990). *Animal Liberation*. New York: Avon Books.
- Steuwer, B. (2021a). Aggregation, Balancing and Respect for the Claims of Individuals. *Utilitas*, 33, 17-34.
- Steuwer, B. (2021b). Limits to Aggregation and Uncertain Rescues. *Utilitas*, 1-14.
- Tadros, V. (2019). Localized Restricted Aggregation. *Oxford Studies in Political Philosophy*, 5, 172-204. doi:10.1093/oso/9780198841425.001.0001
- Temkin, L. S. (2012). *Rethinking the Good*. Oxford: Oxford University Press.
- Tomlin, P. (2012). On Fairness and Claims. *Utilitas*, 24(2), 200-213.
- Tomlin, P. (2017). On Limited Aggregation. *Philosophy & Public Affairs*, 232-260.
- Van Gils, A. (2019). *The Aggregation Problem for Scanlonian Contractualism*. Reading: University of Reading.
- Van Gils, A., & Tomlin, P. (2019). Relevance Rides Again? Aggregation and local relevance. *Oxford Studies in Political Philosophy*, 6, 221-255. doi:10.1093/oso/9780198852636.003.0008
- Voorhoeve, A. (2013). Vaulting Intuition: Temkin's Critique of Transitivity. *Economics and Philosophy*, 29(3), 409-423.
- Voorhoeve, A. (2014). How should we aggregate competing claims? *Ethics*, 125(1), 64-87.
- Voorhoeve, A. (2017). Why One Should Count Only Claims with which One Can Sympathize. *Public Health Ethics*, 10(2), 148-156.

Voorhoeve, A. (2020). Healthy Nails versus Long Lives: An Analysis of a Dutch Priority-Setting Proposal. In N. Eyal, S. A. Hurst, C. Murray, S. A. Schroeder, & D. & Wikler, *Measuring the Global Burden of Disease: Philosophical Dimensions* (pp. 273-292). New York: Oxford University Press.

Wu, P. (2022). Aggregation and Reductio. *Ethics*, 508-525.