

A Defence of the Relevance View in Ethics

A thesis
submitted in fulfilment
of the requirements for the Degree of
Master of Arts (Philosophy)
at the
University of Otago
by
Adam Carter

University of Otago
2020

Abstract of a thesis submitted in fulfilment of the requirements for the Degree of Master of Arts (Philosophy).

A Defence of the Relevance View in Ethics

by

Adam Carter

When different parties have claims on an agent for limited aid, the agent must decide whose claims to satisfy. Individuals' claims have a certain strength depending on how much they have to gain by receiving aid. To determine which party's claims to satisfy, some views say that we should simply weigh the claims of competing parties against each other. Other views say that we should satisfy the set of claims that contains the strongest overall claim and should not do any weighing at all beyond this. The Relevance View accepts that we can weigh competing claims but puts a constraint on when one kind of claim can outweigh another kind of claim. Claims of one kind must be *relevant* to claims of another kind in order to outweigh them. The Relevance View also stipulates that one kind of claim is relevant to another kind of claim if the first claim is sufficiently close in strength to the second claim. This thesis aims to develop the best version of the Relevance View and defend it against the strongest objections. In Chapter 2, I will situate the Relevance View in relation to competing views and determine the best rationale for the view. In Chapter 3, I will defend the Relevance View against various sceptical arguments. In Chapters 4 and 5, I will show that the Relevance View can overcome certain objections from real-world settings when its scope is properly mapped out. In Chapter 6, I will show how the Relevance View can offer guidance when presented with a diverse range of competing claims.

Keywords: Relevance View, Limited Aggregation, Restricted Aggregation, Aggregate Relevant Claims, Weak Aggregation, Distributive Justice

Acknowledgements

I would like to thank my supervisors Andrew Moore and Elizabeth Fenton for their guidance on the project: to Andrew for your encouragement and thoughtful feedback over the year, and to Elizabeth for your helpful and enthusiastic engagement in the final few months of the project.

I am grateful to the University of Otago for receiving the Research Masters Scholarship. I am also grateful to have received the Daniel Taylor Senior Scholarship in Philosophy.

Thank you to all my colleagues in the philosophy department for your support. I am particularly grateful to Joe for your helpful comments, to Jeronimo and Briony for your support throughout the project both intellectually and more broadly, and to João for all the laughs.

I am grateful to my parents, James and Nicole, for all your continued support. To my friend James for your editing advice. Finally, to Masnoon, for all your engagement throughout the entire project, for always looking out for me, and for philosophising with me over the past few precious years.

Table of Contents

Abstract.....	ii
Acknowledgements	iii
Table of Contents	iv
List of Tables	vi
List of Figures.....	vi
Chapter 1 Introduction	1
1.1 Motivating the Relevance View	1
1.2 Method.....	3
1.3 The competing claims framework	3
1.4 Roadmap.....	5
Chapter 2 The Rationale Behind the Relevance View.....	7
2.1 The intuitive problem again.....	7
2.2 The appeal of Pure Aggregation.....	8
2.3 The appeal of Anti-Aggregation.....	9
2.4 The Relevance View.....	10
2.4.1 The Threshold View	11
2.4.2 Maximin.....	12
2.4.3 Scanlon and tie-breaking	12
2.4.4 Conclusion	14
2.5 The Relevance View and deontology.....	14
2.5.1 Two kinds of defeating	15
2.5.2 Numbers scepticism.....	16
2.6 Determining relevance.....	17
2.6.1 Voorhoeve’s test	17
2.6.2 Unanimity objection	18
2.6.3 A tracking problem.....	19
2.7 Weak Aggregation	20
2.7.1 What is Weak Aggregation?.....	20
2.7.2 The Relevance Requirement.....	22
2.7.3 Weak Aggregation and solidarity	23
2.8 Conclusion.....	23
Chapter 3 Sceptical Arguments Against the Relevance View	25
3.1 The <i>World Cup</i> intuition.....	25
3.1.1 What if the fans knew.....	26
3.1.2 Psychological biases	26
3.1.3 Large number scepticism	28
3.1.4 Reliability strategy.....	29
3.2 The <i>Life versus Paraplegia</i> intuition	29
3.2.1 The real impact of certain harms	30
3.2.2 Two domains of morality.....	31
3.3 Hedging our bets.....	32
3.4 Conclusion	33
Chapter 4 Intrapersonal Aggregation and the Unity of Individuals.....	34

4.1	Intrapersonal aggregation	34
4.2	Iterated choices	35
4.3	<i>Ex ante</i> reasoning.....	36
4.3.1	Productive social risk.....	37
4.3.2	Does <i>ex ante</i> reasoning undermine the Relevance View?.....	38
4.4	Two further problems with restricted <i>ex ante</i> reasoning	40
4.4.1	The problem of dependent risk.....	40
4.4.2	The problem of the identified life.....	42
4.5	Conclusion	44
Chapter 5 How Revisionary is the Relevance View?.....		45
5.1	Priorities in healthcare settings.....	45
5.1.1	Defending the treatment of minor ailments: aforementioned resources.....	46
5.1.2	Defending the treatment of minor ailments: new resources	47
5.2	Competing claims and risk	49
5.2.1	A problem for competing claims that involve risk	49
5.2.2	Differentiating between acts and conjunctions of acts	50
5.3	Conclusion.....	52
Chapter 6 Developing the Relevance View.....		53
6.1	Three-way choices and transitivity.....	53
6.1.1	Temkin on transitivity.....	55
6.2	Complex two-way choice cases.....	57
6.2.1	Kamm on irrelevant utilities	58
6.3	Conclusion.....	59
Chapter 7 Conclusion		61
7.1	Chapter summaries	61
7.2	General conclusions.....	63
Appendix A An Alternative Response to the Argument from Risk.....		65
A.1	Differentiating between social risk and individual risk.....	66
Appendix B Epistemic Considerations		70
B.1	Knowledge and <i>ex ante</i> justifications.....	70
B.2	Epistemically irrational? A response to Fleurbaey and Voorhoeve	71
References.....		73

List of Tables

Table 1	An Iterated Choice.....	36
---------	-------------------------	----

List of Figures

Figure 1	Weak Aggregation.....	22
----------	-----------------------	----

Chapter 1

Introduction

We cannot always help everyone. Agents are often faced with situations where they must choose which group to aid. Take the decision of a publicly funded agency to purchase one drug over another. Given the choice between funding a drug that will cure a moderate illness and a drug that will cure a severe illness, the fact that the latter drug will help its target patients by a greater amount seems to count in its favour. However, another factor that seems to matter in making this decision is the *number* of patients that the respective drugs will help. If one drug will help many people in a population while the other drug will help just a few people, this seems to count in favour of the wider-reaching drug.

If you are the potential recipient of one of these drugs, then we can say that you have a *claim* of a certain strength on this agency. Nearly everyone agrees that the strength of a claim matters. What is more controversial is whether the number of claims matter. In this thesis, I will explore a view about the circumstances in which the number of claims matter. I will defend a position whereby the situations where numbers matter are constrained by a notion of relevance. In order for one kind of claim to outweigh another kind of claim, it must be relevant to this claim. One kind of claim is relevant to another kind of claim only when these claims are sufficiently close in strength to each other.

1.1 Motivating the Relevance View

Consider a case where we can either prevent one death or prevent 1,000 people from becoming paraplegics. Assume for the sake of simplicity that there are no differences in our relationship with any of these parties (for example, that the death-threatened person is under our care). Many hold the strong intuition that we should satisfy the claims of the many individuals, each of whom faces a lesser harm than the worst placed individual. One way to explain this is that the aggregation of claims to be saved from paraplegia is enough to outweigh the single claim not to die. Therefore, it seems that the mere number of people with claims can make a difference as to how we should act. However, this general observation is not always intuitive. Consider the following *World Cup* case described by T.M. Scanlon:

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

Many will think that failing to save Jones is wrong and that this is true no matter how many football fans would be frustrated by temporarily stopping the broadcast. The fans are just trying to watch the match whereas Jones is trying to avoid agonising electrocution. One way to explain this intuition is that no number of people with weaker claims can outweigh a person with a much stronger claim. This says something interesting about aggregating competing claims: when sets of claims compete, the way weaker claims work to outweigh stronger claims is not always additive.

Call the position whereby competing claims add up like stones on a set of balance scales Pure Aggregation.¹ Pure Aggregation implies that when a set of weaker claims competes with a set of stronger claims, there is always some number of weaker claims that can outweigh the set of stronger claims. Utilitarianism is a paradigmatic example of Pure Aggregation. Utilitarianism treats claims as directly proportional in strength to the interests that they protect and all claims add to the choice-worthiness of an act.

A rival position on aggregation which can capture our intuition in the *World Cup* case is called Anti-Aggregation. Anti-Aggregation resolves competing claims cases by giving priority to the person with the strongest claim. Because Anti-Aggregation explicitly rejects aggregation, it implies that no number of weaker claims can ever outweigh a stronger claim.

Given the *pro tanto* desirability of identifying a single unified model for how to resolve competing claims, I seek to defend an intermediate position between Pure Aggregation and Anti-Aggregation. This position is called the Relevance View. The Relevance View states that when a set of weaker claims competes with a set of stronger claims because both sets of claims cannot be jointly satisfied, the set of weaker claims can outweigh the set of stronger claims when and only when the competing claims are sufficiently close in strength.² Because paraplegia is (let us suppose) almost as bad as death for its victim, a sufficient number of claims to prevent paraplegia can successfully aggregate against a competing claim to prevent death. So, even though one claim against paraplegia is not enough to defeat a claim to prevent death, 1,000 of these claims certainly seems sufficient. Contrast this with the *World Cup* case. Because the claim of any given football fan in the *World Cup* case is so much weaker than the claim of Jones, the claims of any number of football fans cannot outweigh the stronger claim of Jones. The set of claims of the football fans is not a “relevant” consideration when helping those fans comes at Jones’s expense. This provides a justification for our duty to save Jones.

¹ Temkin (2012, p.25) adopts a broader definition of aggregation: “[A] principle of aggregation will be any function from the relevant elements—whatever those turn out to be—to the assessment of the goodness or the comparative desirability of the item being assessed, say, lives or outcomes.” According to this definition, a theory that simply averages competing parties’ claims to determine which claims to satisfy is aggregative. Aggregation as I understand it is essentially additive.

² This articulation most closely follows Voorhoeve (2014) who calls his view ‘Aggregate Competing Claims’.

1.2 Method

I motivate the Relevance View by drawing conclusions from our intuitions in highly artificial cases. Some will object that nothing informative can be drawn from cases that are so impoverished of the nuance we see in real-world correlates. I have three initial responses to this objection. First, my aims are more explanatory than revisionary. Analysing such cases is important because they isolate morally important phenomena. There is certainly a danger in moving from hypothetical thought experiments to all things considered conclusions in real-world settings. However, my aims are more modest than that. I aim to map out the structure of one component that enters into questions of distributive justice. I use artificial cases merely to map out this position; whether or not we have good reason to think that it is true depends on whether it explains and informs our moral intuitions in real-world cases.

My second response to the objection that my reliance on artificial cases is problematic is that I do not merely rely on these cases to motivate the Relevance View (although they play a central role). I also argue for the view on other sorts of bases. For example, I argue that there is an independently appealing rationale for the Relevance View.

My third response to the objection that my reliance on artificial cases is problematic is that I acknowledge the force of this objection in various points throughout the thesis. For example, I discuss potential challenges for Relevance View from public-policy-like scenarios that involve risk, rather than relying wholly on the artificial set-up of thought experiments under certainty.

Another inevitably artificial element of this topic is that in order to illustrate the Relevance View, I assume a spectrum of claims of increasing strength. For example, I assume that a migraine is worse than a headache, that a severe bodily injury is worse than a migraine, and so forth, up until death. Of course, these different health-related setbacks will affect different people's well-being differently. For example, incurring a bodily injury may be worse for someone than death in some cases. Indeed, death may actually be in someone's interest on certain accounts of the badness of death. However, I use these examples merely to illustrate the view. I do not commit to the view that there is such thing as a universal claim to have one's life saved that is independent of a much richer picture of the actual claimant's circumstances. If the reader disagrees with my ordering, or the strength that I attribute to each kind of claim, they are free to use different examples.

1.3 The competing claims framework

Claims go towards establishing a duty on the part of the agent (Walen 2019, p.8).³ Suppose patients Anne and Barb are each in need of some medicine. However, their doctor, call her Debbie, only has

³ Tadros (2019) asserts that the language of claims commits one to contractualism and prefers to say that duty-grounding facts go towards establishing duties. I think Walen (2019) instead successfully makes use of the language of claims without being constrained by contractualism. Moreover, I choose

one standard dose, and Anne and Barb each require the whole dose to avoid falling ill. Because Anne and Barb are each affected by who Debbie decides to give the medicine to, we can say that Anne and Barb each have a claim on Debbie. In this situation, Anne's claim competes with Barb's claim because their claims cannot both be satisfied. Debbie's duty in this case is determined by the comparative strength of Anne's and Barb's respective claims.

The strength of a person's claim is at least partly determined by the well-being they stand to gain or lose, depending on how the agent acts. Anne's claim reflects her interest in receiving the medicine, which is to avoid falling ill. However, claims can also reflect other factors about the situation (Walen 2019, p.51-3). While I do not commit to any of the following factors, I introduce them merely to highlight the flexibility of adopting the language of claims. One possible factor in the situation we are considering is how well-off Anne and Barb are relative to each other. If both Anne and Barb have an equally strong interest in avoiding illness, but Anne has been burdened with illness her whole life whereas Barb has had a relatively healthy life, then we might think that Anne has a stronger claim than Barb. The relationship between the patient and the agent may also affect the nature of the claim. Suppose for example, that Anne is a registered patient of Debbie but Barb is not; then (so long as this does not violate Debbie's duties as a doctor) Anne's claim may be stronger claim than Barb's claim along this dimension.⁴ A patient may also be able to alter their claim by waiving it. Anne may waive her claim, for example, generating a duty for Debbie to satisfy Barb's claim. Finally, whether the interest is protected by a claim to not be harmed or a claim to aid may make a difference to the strength of one's claim.

Of course, sometimes the parties competing for aid will contain more than one member. Suppose, in another situation, that Anne, Barb, and Claire are each in need of some medicine and that Doctor Debbie only has one standard dose. This time, Barb and Claire only require half a standard dose each to receive the full benefits of the medicine. Debbie can either satisfy Anne's claim to the medicine or instead the claims of both Barb and Claire. In this situation, Barb's and Claire's claims are aligned. Choosing to split the dose between these two people means that both Barb's claim and Claire's claims are met.⁵

The set of competing claims situations is a subset of problems within distributive justice. The competing claims framework does not have anything to say about which distribution is best or ideal. For example, if one party owns a resource and that resource might better serve another party, then transferring this resource might lead to a *better* distribution of resources. However, this might not be

to use the language of claims because it is simpler to grasp, and because many others in the debate do so as well.

⁴ It may even have overriding strength. More on this in §2.5.1. In this sense, "nature" seems a better way to describe claims than "strength" because a claim may have properties that behave differently in different dimensions.

⁵ Determining which set of claims to serve is not the end of the story. For example, the group whose claims were not satisfied may have secondary claims to compensation. See, Walen (2019, p.101).

an instance of competing claims because the first parties' ownership of the resource restricts others having a claim on the resource.

To narrow my focus even more, in the cases I focus on, the agent is always thwarting one group's claims and serving another group's claims. I will not consider cases, for example, where somebody is being used as a means to satisfy another person's claim. Moreover, unless otherwise stated, I also assume that all patients have the same level of well-being and have not waived their claims. This narrower range of cases will have direct implications for how we deal with a broader range of competing claims cases where the morally relevant factors are more numerous.

1.4 Roadmap

The aim of this thesis is to offer a general defence of the Relevance View. My aim is not to show that the Relevance View is the best view all things considered, only that it is a serious contender with a number of merits.

The Relevance View is motivated by the fact that it can capture two important intuitions that seem to pull in opposite directions: from the *World Cup* case, the intuition that aggregation is sometimes morally inappropriate; and from the *Life versus Paraplegia* case, the intuition that aggregation is sometimes morally appropriate. The Relevance View explains why aggregation is morally appropriate in some cases but not in others: we should only aggregate when competing claims are sufficiently close in strength. But without something deeper said in favour of it, the Relevance View does seem to lack a theoretical basis that, I will argue, two of its main rivals—Pure Aggregation and Anti-Aggregation—share. I address this issue in Chapter 2.

Chapter 2 has dual aims. The first is to situate the Relevance View in relation to rival theories and within the tradition of deontological moral theory. The second aim of Chapter 2 is to establish a rationale—a theoretical source of appeal—for the Relevance View.

In Chapter 3, I will consider a number of expressions of scepticism about the intuitions that motivate the Relevance View. Both proponents of Pure Aggregation and proponents of Anti-Aggregation have come up with creative ways to explain why our intuitions in the *World Cup* and *Life versus Paraplegia* cases are mistaken. I will argue that our intuitions are not vulnerable to these arguments. I also entertain another sceptical worry about finding a single model to adjudicate between competing claims.

Chapters 4 and 5 aim examine the Relevance View in various practical settings including iterated choices, situations involving risk, and the provision of publicly funded healthcare. These chapters aim to show that the Relevance View's scope of application in these settings is limited.

In Chapter 4, I will discuss how the Relevance View can address problems to do with iteration and risk. If an agent repeatedly acts in accordance with the Relevance View, it seems that she commits herself to bringing about situations where everybody is worse off than if she had not acted in this way. It would be desirable if we could find a justification for why the Relevance View would not apply in these situations. The next challenge to the Relevance View I consider is when taking a risk is in everyone's *ex ante* interest, but seems to generate a violation of the Relevance View after the risk has taken place. I argue for how best to respond to this challenge.

Chapter 5 builds on the ideas of Chapter 4. One counter-intuitive implication of the Relevance View is that it seems to recommend moving all of our public health resources that currently treat minor ailments to treating major health-burdens. I explore ways in which the truth of the Relevance View might be consistent with our resourcing of treatments for minor ailments. For example, there may be further democratic reasons to keep our practices as is. I also expand my discussion of how Relevance View can deal with risk where not all parties benefit from the risk being taken.

In Chapter 6, I explore recent attempts to extend the Relevance View to cases that are more complex than the simple cases that motivate the view. Up until this point I only use cases where a set of claims of a given strength compete with another set of claims that are either (a) much stronger, (b) slightly stronger, or (c) the same strength. However, a new set of problems emerge when three groups rather than two groups compete for the same resource. This is another important hurdle if the Relevance View is to be able to explain a broad range of cases that might be relevant to real-world settings such as healthcare prioritisation. Chapter 6 concludes the substantive chapters of this thesis.

In Chapter 7, I begin by offering a summary of what I have argued in Chapters 2–6. I then move onto drawing some general conclusions about the scope of my arguments and identifying some areas for future research.

This thesis also includes two short appendices. Appendix A and Appendix B each respond to a major focal point in Chapter 4. Appendix A is on an alternative line of response to the Argument from Risk that I argue is ultimately unsuccessful. Appendix B addresses epistemic considerations that have been debated in relation to *ex ante* reasoning. My reason for separating these discussions out as appendices is to maintain the momentum of Chapter 4.

Chapter 2

The Rationale Behind the Relevance View

In this chapter I will work towards establishing the strongest rationale for the Relevance View. The Relevance View posits that a weaker set of claims has the potential to outweigh a stronger set of claims only if the weaker claims are *relevant* to the stronger claims. The Relevance View is initially motivated by explaining two intuitions that neither Pure Aggregation nor Anti-Aggregation can explain alone. I explain the cases that motivate these intuitions again in §2.1. The Relevance View is often characterised as a middle way between Pure Aggregation and Anti-Aggregation because it can capture both of these intuitions. In §2.2 and §2.3 I expose the rationales behind Pure Aggregation and Anti-Aggregation, respectfully. In §2.4 I distinguish the Relevance View from other related views. In §2.5 I situate the Relevance View within a wider moral tradition. This sets us up to start evaluating a rationale for the Relevance View. In §2.6 I outline and appraise Alex Voorhoeve’s (2014) idea that relevance is essentially related to taking up what he calls the “permissible personal perspective”. In §2.7 I outline and appraise an alternative rationale due to Alec Walen (forthcoming) called Weak Aggregation. I argue that the second of these proposals is the more successful.

2.1 The intuitive problem again

Pure Aggregation implies that there is always some number of weaker claims that can outweigh a set of competing stronger claims. However, this seems incompatible with our intuition in the *World Cup* case (Scanlon 1998, p.235). In the *World Cup* case, Scanlon asks us to imagine a scenario where we can save Jones from agonising electrocution in the broadcast station only by stopping the broadcast and frustrating some large number of football fans. It seems obvious to many that we should stop the broadcast and save Jones, and that this is true no matter how many football fans will be thereby disrupted. Contrary to Pure Aggregation, the *World Cup* case intuition implies that no number of weaker claims to have the broadcast continue can outweigh Jones’s claim to be saved from electrocution.

Anti-Aggregation holds that we commit a kind of moral mistake when we aggregate claims across individuals. Claims just do not “add up” in the way Pure Aggregation assumes they do. Without aggregation, identifying the overall strongest claim is one way to decide which set of competing claims to satisfy. This criterion provides a justification for satisfying Jones’s claim in the *World Cup* case, because in that case Jones has the strongest individual claim. However, many people feel that aggregation remains appropriate in situations where the competing claims are closer together in strength. Consider, for example, *Life versus Paraplegia*: even if death is worse than paraplegia, given the choice between preventing one death and preventing 1,000 cases of paraplegia, many hold the intuition that we ought to satisfy the claims of the potential paraplegics. Pure Aggregation and Anti-

Aggregation resolve in distinct ways cases where claims compete. But can some broader theoretical considerations capture the source of the appeal of these divergent positions?

2.2 The appeal of Pure Aggregation

I follow Nagel (1979) in thinking that the attraction of various theories of distributive justice is the way that they conceptualise equality.⁶ Utilitarianism is the most influential form of Pure Aggregation. Nagel conceptualises utilitarianism's interpretation of equality as follows:

The moral equality of utilitarianism is a kind of majority rule: each person's interests count once, but some may be outweighed by others. It is not really a majority of persons that determines the result, but a majority of interests suitably weighted for intensity. Persons are equal in the sense that each of them is given a 'vote' weighted in proportion to the magnitude of his interests. Although this means that the interests of a minority can sometimes outweigh the interests of a majority, the basic idea is majoritarian because each individual is accorded the same (variable) weight and the outcome is determined by the largest total. (Nagel 1979, p.112)

Other popular theories of distributive justice build on utilitarianism's understanding of equality by also being sensitive to people's level of well-being. One common view is that the worse off someone is, the more that their interests count. To take Nagel's analogy further, each person's vote is not only weighted in proportion to the magnitude of their interests, but it is weighted more if they are worse off. These theories come under the name of either prioritarianism (Parfit 2012) or egalitarianism (Otsuka & Voorhoeve 2009). Prioritarianism determines how poorly someone is doing on an absolute level, whereas egalitarianism determines how poorly someone is doing by comparing them to others. In this sense egalitarianism is a relational account of priority.⁷

By their very nature, all theories of Pure Aggregation are unable to rule that we should always save Jones from electrocution in the *World Cup* case. It does not matter how much we strengthen Jones's claim to reflect how poorly he is doing; some number of weaker claims can always outweigh Jones's claim. (There is always some number of grains of sand that can outweigh a jumbo jet.)

⁶ As I introduce these theoretical sources of appeal, I mean to be neutral between consequentialism and non-consequentialism. For example, Maximin rejects Pure Aggregation and can be represented as a consequentialist position. Maximin could be represented as the position that the best state of affairs is where the worst-off person is as well-off as they can be. Here, everyone else's interests make no difference to the value of the state of affairs.

⁷ One question that can be asked of both of these theories is: how to understand the scope of well-being that is the basis for attributing priority? For example, is priority based on how well someone's life is going on the whole or how they are doing at the time they face the threat? Scanlon (1998, pp.227-8) offers a more flexible idea when he suggests that different aspects of well-being will be relevant in different contexts. For example, suppose an agent can alleviate two out of five months of pain for Anne, or two out of two months of pain for Barb. On a relational understanding of priority, Anne's claim would be strengthened relative to Barb's claim to reflect the fact that Anne faces three additional months of pain regardless of whose claim is ultimately satisfied. So, the fact that Anne had a bad childhood, for example, may not necessarily be a relevant *aspect* of welfare in this case.

Thus, Pure Aggregation can be understood as the position that proper moral concern for all claimants requires treating their claims as adding to or subtracting from the choice-worthiness of an act over alternative acts. As well as justifying intuitions elicited in such cases as *Life versus Paraplegia*, where a number of claims against paraplegia outweigh a single claim against death, Pure Aggregation finds support from the attractive principle that links proper moral concern to the idea that any claim can make an actual difference to the way an agent should act in any given circumstance. If each fan's claim weighs equally in an agent's decision, then enough of these claims should be able to tip the balance of claims in favour of letting the broadcast continue. This contrasts with Anti-Aggregation where agents are guided to act on the basis of just the strongest overall claim. How appealing one finds Pure Aggregation's conceptualisation of equality will influence how difficult it will be to bite the bullet in the *World Cup* case.⁸

2.3 The appeal of Anti-Aggregation

Nagel (1979, p.125) contrasts his own egalitarian view with a "more uncompromising type" of priority for the person with the strongest claim. Nagel argues that this priority can be defended using a form of moral reasoning he calls "pairwise comparison". Nagel believes "that the general form of moral reasoning is to put yourself in other people's shoes" (p.126). He continues:

[This] includes a separate concern for each person, and it is realized by looking at the world from each person's point of view separately and individually, rather than by looking at the world from a single comprehensive point of view. Imaginatively one must split into all the people in the world, rather than turn oneself into a conglomeration of them.

According to Nagel, the most natural way to choose amongst competing claims while staying true to everyone's separate point of view is to "give priority to the most urgent individual claims" (p.127). Suppose that Anne and Barb have competing claims on our aid, and that Anne's claim is stronger than Barb's claim. When one compares these claims, one is compelled to serve Anne's claim. To further illustrate this pairwise method, suppose that Claire also has a claim on us for aid and her claim is aligned with Barb's claim. If Claire's claim is also weaker than Anne's claim, then our pairwise judgement is unchanged: Anne's claim presents itself as the most important when compared to either Barb's or Claire's claim. Staying true to the pairwise method compels us to serve Anne's claim. This

⁸ One further reason Pure Aggregation seems attractive in the domain of competing claims we are considering is that claims merely vary in strength along the same dimension. Contrast this with situations where claims are of different kinds. For example, if Elma owns a resource and this resource could make a great number of people happy for a day, we might think that Elma's property claim to this resource prevents (or protects) other peoples' claims from outweighing her claim. This makes sense because Elma's claim is of a different kind to the claims of the would-be happy people. However, in the cases we are considering, all claims fall on the same continuum of strength, they all related to something like health-related well-being, and so we might presume that all claims in these settings should have equal marginal importance.

is one explanation of our intuitions in the *World Cup* case. When comparing Jones's claim to the claim of each football fan taken alone, serving Jones's claim is far more important.

Anti-Aggregation rejects the intuition in the *Life versus Paraplegia* case, and would instead recommend serving the single claim against death rather than the 1,000 claims against paraplegia on the grounds that this single claim is the strongest claim. However, this approach appears to offer no solution when the competing claims are equal in strength. If Anne's claim is equally as strong as each of Barb and Claire's claims, nobody's claim is the most urgent. One common suggestion in this situation is that all an agent can do is offer equal chances through a fair procedure, such as a coin toss (Taurek 1977).

As Nagel himself recognises, these two different ways of resolving competing claim conflicts rely on different and credible justifications. Pure Aggregation treats proper moral concern for each individual as taking their claims to "add up" together with the claims of other individuals. Anti-Aggregation treats proper moral concern as staying "closer to the points of view of the individuals considered" (p.123).

2.4 The Relevance View

The Relevance View is a middle way between Pure Aggregation and Anti-Aggregation. According to the Relevance View, aggregation is morally appropriate in some cases and not in others. Aggregation is appropriate only when competing claims are 'relevant' to one another where relevance is determined by whether two competing claims are 'close enough' in strength.⁹ In the *Life Versus Paralysis* case, the two kinds of competing claims are significant relative to each other in that they are "close enough" in strength, and so aggregation is appropriate. This is to say that some number of claims against paralysis could defeat a competing claim against death. Things are different when we consider the *World Cup* case. Jones's claim to be rescued from electrocution is much stronger than the claim of any given football fan to watch the match on TV. The difference in what is at stake means that these competing claims are not relevant to each other. Therefore, no number of these weaker claims can defeat Jones's stronger claim.

I will now work towards uncovering the theoretical appeal of the Relevance View, to establish the view as a plausible contender alongside Pure Aggregation and Anti-Aggregation. My first task will be to situate the Relevance View by differentiating it from and highlighting its advantages over related views.

⁹ Voorhoeve (2014, p.65 fn. 2) offers a list of philosophers who have defended a view like this.

2.4.1 The Threshold View

One way in which consequentialists have tried to accommodate our intuitions in the *World Cup* and *Life versus Paraplegia* cases is by introducing the idea of value superiority. Mill (1961, II 6) famously makes the distinction between higher and lower pleasures, such that no number of lower pleasures can outweigh a higher pleasure. Using this idea in the current setting, when claims cross some threshold of strength, no number of claims below this threshold can outweigh the claims above this threshold. The Threshold View could say that Jones's claim and the claims of the football fans lie on either side of this threshold, whereas the claims of the potential paraplegics and the claim of the person facing death each lie above this threshold.

However, the Threshold View seems to lead to counter-intuitive results when the claims being compared fall just on either side of the threshold. For example, if this threshold is equivalent to a claim against having two broken limbs, then the Threshold View would imply that no number of claims to have a single limb spared can outweigh even a single claim to have three limbs spared. However, these claims seem close enough in strength such that enough of one claim can outweigh the other.

By contrast, the Relevance View holds that whether one claim can outweigh another claim depends on the strength of these claims *relative* to each other. For example, a weak claim may be relevant to a moderate claim, a moderate claim may be relevant to a strong claim, and yet a weak claim may not be relevant to a strong claim.

A related idea is that morality is just not concerned with trivial claims such as the claim to have the football broadcast continue. For example, if duties of beneficence are concerned only with others' life-projects, then it is not obvious that mild doses of frustration are objects of beneficence at all. Another way of putting it is that the interests of the football fans in the broadcast continuing do not ground beneficence claims on an agent.¹⁰ My first objection to this idea is that finding the point at which people's interests are no longer the objects of beneficent concern will face similar problems to those just outlined, where claims falling on either side of this threshold will seem relevant to each other. My second objection is that while it is true that trivial interests do not typically ground beneficence duties (i.e., are the objects of claims), this is not because they cannot ground such duties. Instead, it is because the beneficence duties they would ground would typically be too demanding. If I were to go out of my way every time that I could fulfil another's very small interest, I plausibly would be robbed of the freedom to live my own life. Finally, if the electrical equipment had not fallen on Jones, then turning off the broadcast for no other good reason would definitely be wrong. This would be a case where 'trivial' interests do ground duties of beneficence.

¹⁰ Hirose (2015, p.46) raises a similar question but it is not clear how his use of trivial differs from relevance. See also Bramble (2019) where he argues that only impacts on lifetime well-being have normative significance.

I conclude that the Relevance View has distinct advantages over alternative views that aim to explain our intuitions in the motivating cases using absolute thresholds.

2.4.2 Maximin

In this section I contrast the Relevance View with an Anti-Aggregationist rival called Maximin. Maximin is the principle that we should distribute resources so as to make the worst off people the least worst off that they can be.¹¹ The idea is to maximise the well-being of the people with the least amount (minimum) of well-being. Thus described, Maximin is a distinctly Anti-Aggregationist view. However, another contrasting feature of Maximin and the Relevance View is the scope of the claims in consideration. Suppose that some football fans in the *World Cup* case have had particularly impoverished lives whereas Jones has lived a comparably privileged life. Maximin would recommend letting the broadcast continue in this case because the worst off individuals can be made ever so slightly better off.

However, the Relevance View may still recommend saving Jones, who has the strongest claim in this much more narrow domain of competing claims. Note also that the Relevance View may still conceive of claims as reflecting how poorly someone is doing as well as how much they have to gain by the agent acting in their favour. If some number of football fans have lives that are going much worse than Jones's life, then the claims of the football fans may be stronger than if they were doing as well as Jones. Still, so long as the claim of any given football fan is sufficiently weaker than Jones's claim, the Relevance View would recommend saving Jones. Therefore, at least in the context of competing claims situations, the Relevance View is superior to Maximin.

2.4.3 Scanlon and tie-breaking

Scanlon (1998) presents his moral theory as distinctly anti-aggregative in nature.¹² However, Scanlon argues that we may be justified in letting numbers break ties. Suppose, for example, that Doctor Debbie only has Anne and Barb as patients, and they have competing claims on Debbie for the single dose of medicine. In this case, Debbie has a permission to help either individual. However, before Debbie acts, Claire is admitted to the hospital and is also in need of the same medicine. Barb and Claire only require half a dose of medicine while Anne requires a full dose, so Debbie can either give the dose to Anne or split it between Barb and Claire. Scanlon (1998, p.232) argues that in cases such as this, it would be disrespectful to Barb and Claire to save Anne; it would be as though the addition of either of their claims to the competition made no difference and counted for nothing. According to

¹¹ Hirose (2015, p.28). Maximin is most famously expressed in the writings of Rawls (1971).

¹² Scanlon's idea that individuals alone can reasonably reject principles is the foundation of the theory. One circumstance of reasonable rejection concerns how the most burdened person fares under this or that principle. An individual can reject a principle if they would be better off under another principle and nobody else would be more worse off than them (p.230).

Scanlon's justification, because numbers can break ties, Debbie may be required to split the dose.¹³ This tie-breaking argument is anti-aggregative because the disrespect we show Barb is independent of the disrespect we show Claire.

It would be a mistake to ignore Scanlon's own contribution to the development of the Relevance View. Scanlon (1998, p.238-41) explores a way of extending the tie-breaking argument to work over a broader range of cases. Suppose that Barb and Claire each have a claim that is only slightly weaker than Anne's claim. Scanlon feels inclined to say that the agent may still be permitted to save Barb and Claire. Saving Anne would ignore the fact that the claims of Barb and Claire are almost as strong as the claim of Anne. Scanlon's (1998, p.239) tolerance of numbers is thus restricted by a notion of *relevance*: the number of claimants can break ties only if the competing claims, taken individually, are close enough in significance to each other.

However, the major problem with Scanlon's argument is that there is a certain arbitrariness in allowing tie-breaking and not aggregation. After all, tie-breaking already acknowledges that claims are commensurable (Liao 2008, p.445), and so goes beyond looking at claims in isolation as the pairwise method, for example, would have us do. To break ties, claims must be matched against each other. In my view, tie-breaking is best thought of as a heuristic for aggregation: it allows us to decide on the basis of what is left over when competing claims cancel, and this is easier than determining the aggregate strength of competing claims. Rather than attributing to each claim some determinate strength and weighing these claims against each other like we would with rocks on a set of balance scales, tie-breaking allows us to count the mere number of claims on each side to determine a verdict without determining the strength of any given claim (so long as they are close-enough in strength).

Verdicts based on tie-breaking with relevance would come apart from the Relevance View because the tie-breaking argument counts the mere number of claimants, and is not sensitive to the strength of the claims as long as they are relevant. When large numbers are involved, an error will accrue if we rely only on tie-breaking because tie-breaking is not being sensitive to the marginal strength of each claim. Suppose, for example, that 100 claims against death compete with 101 claims against paraplegia and claims against death are relevant to claims against paraplegia. Scanlon's method would have us choose to satisfy the claims of the largest number of claimants, the people facing paraplegia. On the other hand, the Relevance View would have us choose to prevent the deaths. Even if the competing claims were closer together in strength, ignoring the significance of these differences on the individuals affected seems unnecessary.

¹³ Hirose (2015) points out that this reasoning is fully consistent with a weighted lottery. Suppose that Anne is given a blue marble and Barb a red marble in a lottery. Claire's addition to the group could contribute another red marble to the lottery rather than determining that Doctor Debbie now ought to serve both her claim and Barb's claim. This way Anne still has a chance of having her claim served, albeit only a 1 in 3 chance.

2.4.4 Conclusion

I conclude that the Relevance View is different from and has advantages over the Threshold View, Maximin, and Scanlon's tie-breaking argument. The Relevance View determines relevance on the basis of the comparative strength of competing claims rather than on the basis of whether these claims fall on either side of a threshold as the Threshold View does. The kind of concern that the Relevance View shows for the person with the strongest claim is more narrowly construed than in Maximin, and unlike Maximin, the Relevance View is not Anti-Aggregative. The Relevance View also explicitly accepts aggregation in certain circumstances, unlike Scanlon who tries to arrive at similar conclusions using some variety of "proto-aggregation". Now that I have differentiated the Relevance View from related views, I will work towards deepening my theoretical characterisation of the view.

2.5 The Relevance View and deontology

In this section I advance a deontological way of understanding the Relevance View. Deontology is a particularly illuminating way of understanding the Relevance View, although I am agnostic as to whether the Relevance View can be defended within consequentialism (see, for example, Temkin 2012). I understand deontology simply as any view that constrains how agents may (permissibly) bring about certain outcomes. I focus on deontology rather than consequentialism for several related reasons.

First, the notion of relevance can be considered just one of many constraints on maximising the good that deontologists have explored. Some argue that when one acts in such a way as to ignore relevance, there is a wrongful kind of *disrespect* shown to the person with the stronger claim (e.g., Kamm 2007). This idea also fits with another contrast to consequentialism: that there is more than one response towards things that have fundamental moral importance. Whereas consequentialists claim that the only appropriate response to the good is that of *promotion*, many non-consequentialists claim that there are other appropriate responses, such as respect.¹⁴ I will draw on this idea at multiple times throughout this thesis.

The second reason I explore the Relevance View within deontology is that introducing relevance seems to give rise to certain intransitive relations. Transitivity is where if A bears a relation to B and B bears the same relation to C, then A bears this relation to C. For example, if A is taller than B, and B is taller than C, then we can assume that A is also taller than C. By contrast intransitivity is where the consequent is denied. For example, where A is better than B and B is better than C and yet A is not better than C we can conclude that the "better than" relation is not transitive. Intransitivity of the 'better than' relation is something consequentialists are reluctant to embrace but the intransitive 'ought to, rather than' relation poses no particular challenges to the non-consequentialist (Kamm 2007;

¹⁴ See, for example, Swanton (2003).

Temkin 2012). Thus deontology seems more accommodating of intransitivity than consequentialism. I will explore questions of transitivity in Chapter 6.

In Chapters 4 & 5 I show that the notion of waiving a claim can make sense of how an agent should act under conditions of risk. Without pre-empting that argument here, I will note that the notion of waiving one's claim is particularly suited to deontology. One of the key tenets of deontology is respecting a person's powers over the domain of their own welfare. Consequentialism on the other hand, with its focus on interests, doesn't offer individuals the same discretion over the duties of a beneficent agent.¹⁵

A fourth reason for favouring deontology in this setting is that relevance might be just one species of a more general *defeating* relation within deontology (Tadros 2019). I will explore this idea now because it demystifies the notion of relevance that is the key insight behind the Relevance View.

2.5.1 Two kinds of defeating

The Relevance View seems to make use of two possible ways in which claims can interact, or more specifically, ways that one claim can defeat another (Tadros 2019). First, one set of claims can defeat another set of claims simply because it outweighs this set of claims. In the *Life versus Paraplegia* case, this is one very plausible explanation of the agent's duty to serve the claims of the many potential paraplegics, rather than the claim of the single person facing death. Because the many claims against paraplegia outweigh the single claim against death, the claims against paraplegia prevent the claim against death from establishing a duty.

By contrast, it seems incorrect to say that Jones's claim outweighs the aggregate claims of the fans. If a weighing relation were involved, then we would expect that the claims of some number of fans would outweigh Jones's claim. Instead, as the Relevance View suggests, the claims of the fans are not relevant when satisfying these would mean leaving Jones to suffer painful electrocution. We can say instead that Jones's claim *disables* the force of any given fan's claim from its potential to outweigh Jones's claim.

Tadros (2019) offers another example of this relation at play in deontology. The example is between the duty of a doctor and the consent of their patient. Even if it would be strictly better for a patient that their doctor operate on them without their consent, if the patient does not consent and they have the capacity to do so, then the absence of their consent *disables* any duty to promote the good that the doctor might have in this case.¹⁶ In this case, without consent, the prospect of promoting someone's

¹⁵ Of course, consequentialists often appeal to 'subjective' interests or preferences. While this lets people's own preference determine what is good for them, it does not give them the *power* to withdraw this preference at their will.

¹⁶ Cullity (2018) calls a similar phenomenon context-undermining: for example, patient A has control of a domain of his well-being such that the absence of his consent disables any opposing duties of

welfare does not translate into a duty to promote their welfare. A similar phenomenon explains why the severity of Jones's claim prevents the claims of the football fans from grounding a duty on our part that their claims be served. The analogy works well because just as the size of the benefit to the patient does not seem to make a difference to whether the doctor has a duty to perform the surgery without consent, the number of fans does not seem to make a difference to whether the agent should save Jones.

The constraint of respect is important in explaining both cases. The doctor disrespects the patient by assuming that a large enough benefit outweighs their consent. Likewise, we disrespect Jones in the *World Cup* case when we act as though very small annoyances, even in great numbers, are the kinds of things that can outweigh his claim. This constraint on maximising the good is a common theme in the deontological moral tradition.

2.5.2 Numbers scepticism

Some people think that any kind of aggregation violates a core tenet of deontology: the Separateness of Persons. According to proponents of the Separateness of Persons, when Pure Aggregation adds claims together to determine which set of claims to serve, it fails to treat each individual (and their claim) as separate from any other. One way of pressing the Separateness of Persons complaint is that trade-offs between the interests of competing individuals (or groups) do not resolve themselves in the same manner, or as easily, as does the trade-off of interests within a single person. When we elect to have a surgery, for example, we often set back our short-term interests for the sake of satisfying a more substantial long-term interest (such as ensuring mobility for an extended period of time). However, when we impose a burden on one person so that a different person can receive a more substantial benefit, that is no compensation for the person on whom the burden was imposed.

However, the Relevance View is silent on how to address *intrapersonal* trade-offs and so it is not obvious how the Relevance View violates the Separateness of Persons and therefore a central deontological constraint. The Relevance View only applies in situations that involve competing claims belonging to different individuals. Moreover, in the cases we have been considering we are not imposing a burden on someone to benefit another. Nowhere does the view imply that satisfying one set of claims compensates for the claims unsatisfied, if the agent could satisfy both sets of claims then they would. Instead, the groups involved have competing claims to our aid; we are not harming one group to benefit another. To the extent that harming one person in order to benefit another is permissible, this must be because the claim of the person to receive this benefit is significantly stronger than the claim of the person being harmed.

beneficence. Also, there are some difficult cases where consent is difficult to obtain and so this disabling relation may not hold.

Walen (2019) notes that many of deontologist's favourite puzzles rely on interpersonal aggregation for their intelligibility. In the standard trolley problem, for example, a trolley is heading towards five people. A bystander can either let the five die or pull a lever, diverting the trolley onto a side track with one person. Many deontologists endorse the principle that killing is worse than letting die. As Walen (2019, pp.51-2) correctly observes, if deontologists attribute a stronger claim to the person on the side track than to any individual in the way of the trolley, then if diverting the trolley is morally required, this must be entirely because of the numbers. Walen (2019, p.53) also reminds us that when deontologists claim that killing is worse than letting die, they are not saying that it is worse for the one who dies. They are merely saying (to put it abstractly) that killing should weigh more negatively in determining an agent's action than letting die does.

I conclude this section by restating that the Relevance View is firmly at home within deontology. The view readily makes use of the distinction between outweighing and disabling where the disabling relation is plausibly understood as a constraint on maximising the good. Finally, the Relevance View violates no definitive deontological restriction.

2.6 Determining relevance

Even if the Relevance View captures our judgements in the *World Cup* case and *Life versus Paraplegia* case, and relevance is just an instance of the disabling relationship we see elsewhere in ethics, we might think that the Relevance View requires a principled way of determining when one claim is relevant to another. One reason that determining a test is particularly important is because people's intuitions will diverge widely. Pure Aggregation takes this to an extreme by taking the position that all competing harms are relevant to each other. In this section, I explore Voorhoeve's attractive method for determining when competing claims are relevant to each other.

2.6.1 Voorhoeve's test

Suppose that two groups, A and B, have competing claims on some resource, and the claimants within each group have claims of the same strength. This condition holds in the *World Cup* case where the football fans have identical claims and these claims compete with Jones's claim. This condition also holds in the *Life versus Paraplegia* case, where the claims of the many people facing the threat of paraplegia compete with the claim of one person facing the threat of death. Voorhoeve believes that relevance can be determined by analysing a competition between a single claimant from Group A and a single claimant from Group B. Call this a one-one competition and call a representative from Group A, A, and a representative from Group B, B. In a one-one competition, we first need to ask whether or not A should prefer that B's claim be satisfied rather than his own claim out of minimal moral concern for B. If A may prefer that his own claim be satisfied rather than B's claim, then claims of the same strength of A's claim are relevant to claims of the same strength of B's claim. On the other hand, if A should prefer that B's claim is satisfied rather than his own, then claims of the same strength of A's

claim are NOT relevant to claims of the same strength of B's claim. According to Voorhoeve, the legitimacy of A's preferences in these cases is guided by "minimal moral concern" for B. Minimal moral concern involves *some* preference for the satisfaction of one's own claim: A may legitimately prefer the satisfaction of his own claim as long as B's claim is not "too much" larger than A's claim.

To illustrate how this works, consider a case in which members of Group A face the prospect of a headache, and members of Group B face the prospect of an early death. In a one-one competition, minimal moral concern seems to require that A prefer the satisfaction of B's claim over his own. Therefore, claims to have one's headache prevented are not relevant to claims to have one's life saved. Consider a different case in which members of Group A now face the prospect of losing a limb. In a one-one competition, minimal moral concern does not seem to require that A prefer the satisfaction of B's claim over his own. A may retain his preference that his own claim be satisfied over B's claim. Therefore, claims to have a limb saved are relevant to claims that one's life be saved.

Voorhoeve's test for relevance seems to match our intuitions in the *World Cup* and *Life versus Paraplegia* cases. If I were the only one watching the football match broadcast I should prefer that the broadcast be stopped for the sake of Jones. On the other hand, if a trolley is either going to paralyse me or kill another, it would be appropriate for me to prefer that I be saved. But what other reasons does Voorhoeve have for favouring his test? Voorhoeve (2014) believes that "a form of unanimity emerges" (p.74) when everyone with a weaker claim realises that, in a one-one competition, they should prefer that the competing claimant's claim should be served. An individual can hardly complain that their claim is not satisfied in a one-many competition—so this reasoning goes—if they would be required to prefer that their competitor's claim be satisfied in a one-one competition. I will now explore two objections to Voorhoeve's test.

2.6.2 Unanimity objection

My first objection to Voorhoeve's test for relevance is that it is not clear why this one-one relation should necessarily translate into any many-one relation. Tadros (2019, p.177) explains that when I am the only one who can provide someone with a great benefit at a minor cost to myself, I plausibly have a duty to do so. However, when providing the same benefit to someone would require many people to incur an identical minor cost, it is less clear that each of us has a duty to do so. Tadros gives the following example:

[S]uppose I am watching television and I see a person struggling to get his shopping home. I might be required to stop watching and help. It does not follow that any number of people who are watching television must all stop to help if that is the only way to relieve the person's burden. (Tadros 2019, p.177)

Whether I am morally required to benefit someone may depend on whether or not other people would be similarly burdened in order to provide this benefit. No doubt, Tadros uses an odd counterexample, but I think it successfully diagnoses the flaw in Voorhoeve's test. Voorhoeve needs to say why the

unanimity that emerges from each person considering their morally informed preference in a one-one competition has anything to do with how the agent should choose in a many-one competition.

2.6.3 A tracking problem

Even granted the conclusion of the previous section, patients' moral preferences in one-one competitions and an agent's duties in one-many competitions rely on the relative strength of competing claims. If the concept of "minimal moral concern" is correct, a claimant's preference that their competitor's claim be satisfied depends on how much stronger their competitor's claim is than their own claim. Similarly, if the Relevance View is correct, an agent's duty to satisfy one set of claims over another set of claims depends on how much stronger claims of one kind are than claims of another kind. Therefore, Voorhoeve might have simply identified a helpful correlation: we can use one-one competitions to guide us in claims competitions.

However, Tadros (2019) argues that there is precedent in ethics for thinking that what is acceptable for a claimant is different from what is acceptable for an agent. When A endorses, out of minimal moral concern for B, satisfying B's claim rather than his own, he is making a kind of personal sacrifice. It is true that what A is required to sacrifice for the sake of B correlates with the burden we are permitted to impose on A for the sake of B. However, in the rescue cases we have been considering, the situation looks very different from the perspective of the agent. Failing to serve one set of claims is a mere side effect of serving the competing set of claims. In deontology, the distinction between imposing a burden on someone as a means to benefiting someone else and letting a burden befall someone as a side effect of helping someone else, is not trivial. Many theorists assert that we are permitted to allow a greater burden to befall someone as a side effect of our actions than if that burden is a means to achieving our action. Consider the common example of a trolley threatening the life of many people. It is generally considered harder to justify pushing someone off a bridge in order to stop the trolley from killing many people than diverting a trolley onto side track where it will kill one other. If this distinction is justified, then Voorhoeve's test does not track the cases we are interested in.¹⁷ The test may apply to cases where we must harm one group in order to benefit another. A fanciful example of this may be drawing a small amount of blood from the majority of the population in order to benefit a minority via transfusion. By contrast, the cases we are considering such as the World Cup Case do not involve using the fans as a means of benefiting Jones.

¹⁷ This relates to a point raised by Halstead (2016). He worries that Voorhoeve's test raises a problem for non-agential patients (e.g., non-human animals). One's capacity to judge whether one has a duty relies on one's status as a moral agent. If Voorhoeve's test is limited to moral agents, then it leaves an odd asymmetry between moral agents and non-agents who still have the status of moral patients. In later developments of the test, Voorhoeve (2015) stipulates that preferences that one's own claim be satisfied can be determined by the agent taking up the perspective of the patient facing the threat. However, the idea of taking up the perspective of a cow to determine whether it ought to prefer that another cow's harm be averted stretches this idea beyond plausibility.

I conclude that Voorhoeve's test to determine when relevance obtains is unsuccessful. But with no other tests for determining relevance on the radar, is this a problem for the Relevance View? Perhaps not. While determining when one set of claims is relevant to another is no easy task, our judgements at the extremes are clear enough. By way of example, the competing claims in the *World Cup* case are not relevant, while the claim against paraplegia is relevant to death. There are many problems in ethics where it is unclear where one consideration overtakes another and yet we do not feel that we need to design a test. For example, at what point does acting in a child's best interests override their parents' wishes? It does not seem that a simple test will help us decide the appropriate threshold.

Voorhoeve's test served the functions of both test and rationale. It attempts to give a way of determining when one set of claims is relevant to another. It also attempts to explain the Relevance View by showing that Pure Aggregation is appropriate in some cases, and Anti-Aggregation is appropriate in others. I have given up on the of these first functions but the equivalent of Voorhoeve's second function remains important if the Relevance View is to stand its ground against Pure Aggregation and Anti-Aggregation. Pure-Aggregation and Anti-Aggregation have an underlying rationale, what about the Relevance View?

2.7 Weak Aggregation

The Relevance View is the position that when claims compete, we should aggregate only when these competing claims are relevant to each other. Relevance is determined by comparing the significance of individual competing claims, but the judgement of relevance is unlikely to be determined by some test. In this section, I defend another rationale for the Relevance View called Weak Aggregation (Walen forthcoming).

2.7.1 What is Weak Aggregation?

Earlier, I rejected the Threshold View in consequentialism because our judgements about welfare do not seem accommodating to sharp discontinuities in value. However, another kind of sharp discontinuity is present within the Relevance View as described up until this point. A weaker claim is either relevant to a stronger claim, in which case it has full aggregative strength, or it is not relevant, in which case no number of claims of the same strength can defeat the stronger claim. One might dislike this idea that weaker claims have full aggregative strength against a stronger claim up until a point, after which all this strength is lost. Unlike the case of receiving consent, where you either have it or you do not, duties in the context of competing claims do not seem accommodating of sharp boundaries of this sort. Alec Walen speculates that the Relevance View might be part of a larger, profoundly contextual picture of aggregation he calls *Weak Aggregation*:

Weak aggregation treats Limited aggregation [i.e., the Relevance View] as the limit of a more general tendency for weaker claims to aggregate more and more weakly against stronger claims as the gap between them grows. That seems more plausible than the thought that there is a sharp line between those cases in which weak claims add up in the

balance at full strength and those cases in which the ability to meaningfully aggregate suddenly disappears. (Walen forthcoming, p.9)

As a burden is dispersed among individuals, the claims that protect individuals against this burden get weaker. Suppose two people have competing claims to be free of four weeks of pain. Saving either claim seems permissible. Next, consider a case where one person has a claim to be free of four weeks of pain and their claim competes with two other people's claim to each be free of two weeks of pain. In this case, Weak Aggregation implies we should satisfy the claim of the single person who faces four weeks of pain. Finally, consider a case where two people each have a claim to be free of two weeks of pain and their claims compete with four other people's claims each to be free of one week of pain. Weak Aggregation implies we should satisfy the claims of the two individuals who each face two weeks of pain.

We could keep dividing this burden of four weeks of pain equally among individuals and the combined claims that this burden would ground would aggregate more and more weakly against a claimant who bears the burden alone. Eventually, no number of the individually weaker claims could outweigh the burden that is fully concentrated on one person.¹⁸

Weak Aggregation rejects transitivity between the strength of claims. Walen (forthcoming, p.12) illustrates this as follows. Suppose it takes ten weak claims to defeat one moderate claim, and ten moderate claims to defeat one strong claim. If aggregative strength were transitive, we would expect 100 weak claims to defeat one strong claim. However, if Weak Aggregation is correct and the aggregative strength of claims is intransitive, then we could expect that it would take something closer to 1,000 weak claims to defeat one strong claim. Because of this intransitivity, there is no space where claims can be given a universal weighting. This contrasts with the original formulation of the Relevance View, where as long as claims are relevant to one another, they interact in a uniform way that reflects their strength.

¹⁸ Temkin (2012, Chapter 3) discusses a related idea under what he calls Disperse Additional Burdens. However, Temkin seems committed to consequentialism and relatedly seems to be defining aggregation as a function determining the desirability of outcomes.

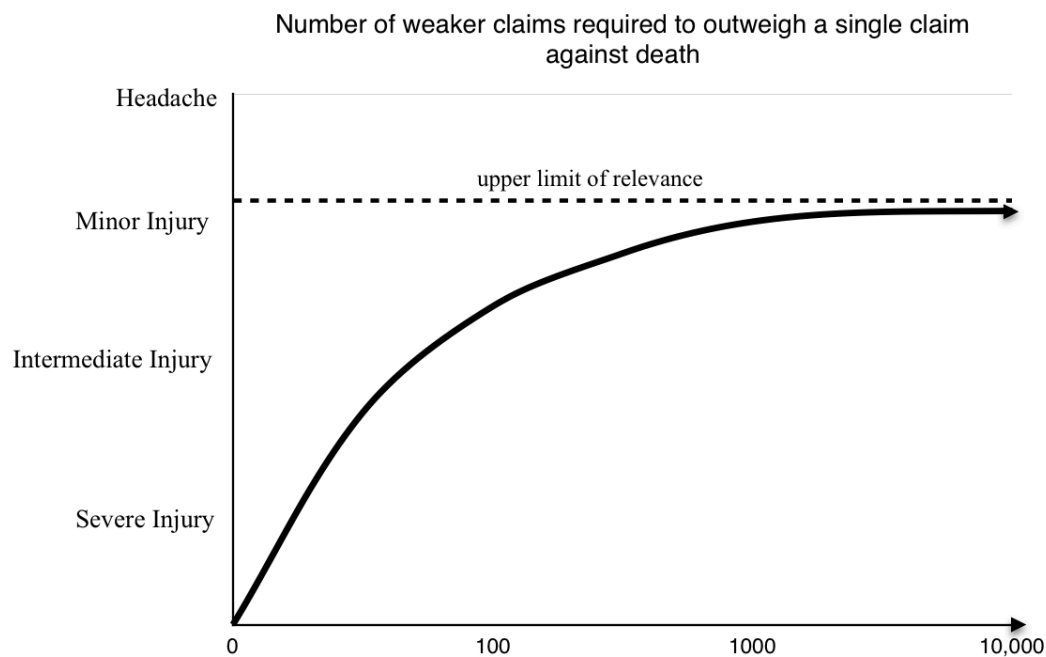


Figure 1 Weak Aggregation

Figure 1.0 is a possible representation of the number of weaker claims required to outweigh a single claim against death. A new graph would have to be drawn if the stronger claim were something other than death. The strength of the weaker claim relative to the stronger claim is represented on the vertical scale. The number of weaker claims required to outweigh a stronger claim is represented on the horizontal scale. According to Weak Aggregation, as the claims get weaker, they diminish in their potential to outweigh the stronger claim until they can no longer do so. This can be represented as a curve that approaches an asymptote. As the strength of a claim approaches the asymptote, the number of these claims required to outweigh the stronger claim reaches infinity. However, above the asymptote, no number of weaker claims can outweigh the stronger claim. At this point, weaker claims can be regarded as no longer relevant to the stronger claim.¹⁹

2.7.2 The Relevance Requirement

Whereas Voorhoeve's proposal was very much a hybrid between Pure Aggregation and Anti-Aggregation, Weak Aggregation revises the Relevance View itself by developing within it a more complex account of the diminishing power of weaker claims to aggregate against stronger claims.

It is not immediately clear where the disabling relation enters this picture. The idea of the disabling relation was a nice way of understanding why continuing with the broadcast in the *World Cup* case is wrong: the claims of the football fans are prevented from defeating Jones's claim because the presence

¹⁹ In many contexts, something like diminishing marginal utility is seen as a bug; in this context, diminishing marginal aggregative strength is a feature. In a related discussion of limited aggregation and risk, Lee-Stronach (2018) discusses the advantages of using bounded values to represent lexical priority over other approaches that use infinite values.

of Jones's claim disables the aggregate force of these weaker claims. However, according to Weak Aggregation (and as indicated on the graph above), when a set of weaker claims compete with a much stronger claim, these weaker claims just lose their power to outweigh the stronger claim as they get weaker. One implication of this story about the diminishing power of weaker claims is that it does not seem that we need to make use of the disabling relation at all. The language we are using is just about weighing. However, the disabling relation could explain something absent rather than something present about how claims interact. It highlights the fact that even though weak claims might ground duties in some contexts, in contexts where they compete with much stronger claims, they do not behave in this way. When we say that the stronger claim disables any number of weaker claims from grounding a duty, this highlights the moral error in thinking that these weaker claims aggregate in a way that relevant claims would aggregate.²⁰

Rather than abandoning the idea of relevance altogether, I suggest we call the phenomena whereby one kind of claim cannot outweigh another kind of claim the 'relevance requirement'. The claims of the football fans do not meet the relevance requirement, so they are unable to conjure up enough strength in the aggregate to defeat Jones's claim. If we satisfy the claims of the fans we violate the relevance requirement.

2.7.3 Weak Aggregation and solidarity

Part of the rationale for Weak Aggregation is that it expresses a general tendency for claims to lose aggregative strength as the burden that grounds them is distributed. I think one reason that this tendency is morally important is that it expresses a kind of solidarity. As a burden is shared by more and more people, people should be more willing to bear the brunt of this burden: "It is not fair that you should have to bear the whole burden yourself. We will each take some weight off your shoulders even if the total burden now increases." Eventually we say that nobody should have to bear a burden when other people can instead bear a relatively insignificant burden.

2.8 Conclusion

Let us conclude this chapter. The Relevance View is at home within deontological moral theory and is plausibly just one instance of the deeply contextual nature of that moral theory's account of moral duties. However, sharp boundaries to do with facts about relevance are not supported by a Voorhoeve-style test, and nor is it desirable to set any sharp relevance-irrelevance boundary. While the Relevance View does not tell a simple story about aggregation as its rivals Pure Aggregation and Anti-

²⁰ The mistake we make when we assume that very weak claims can outweigh much stronger claims may be understood as a kind of "benefit fetishism" (Draper 2002, p.208). We forget that potential benefits belong to individuals and take their normative force from claims. Disrespect in this setting comes from violating the relevance requirement through expressing a kind of "benefit fetishism".

Aggregation do, it does avoid certain unattractive implications of both these extreme views. I have also argued that the Relevance View has its own rationale which expresses the value of solidarity.

The Relevance View is attractive because it reconciles two intuitions that pull in opposite directions. Our intuition in the *World Cup* case pulls us towards Anti-Aggregation and our intuition in the *Life versus Paraplegia* case pulls us towards Pure Aggregation. However, in virtue of reconciling these intuitions, the Relevance View opens itself up to criticism from both Pure Aggregation and Anti-Aggregation. In the next chapter, I explore various attempts to discredit our intuitions in the motivating cases.

Chapter 3

Sceptical Arguments Against the Relevance View

The Relevance View diverges in character from both Pure Aggregation and Anti-Aggregation and so is vulnerable to sceptical arguments from proponents of both views. Both Pure Aggregation and Anti-Aggregation have a *prima face* simplicity that the Relevance View lacks: Pure Aggregation is the view that aggregation is always appropriate, Anti-Aggregation is the view that aggregation is never appropriate, and the Relevance View is the view that aggregation is sometimes appropriate. More specifically, the Relevance View states that aggregation is appropriate when and only when competing claims are relevant to each other (where relevance is determined by being close enough in strength). The simplicity of Pure Aggregation and Anti-Aggregation is at least one *pro tanto* reason to favour them over the Relevance View.

However, the Relevance View is built on the conviction that each of these views is unable to explain two important intuitions. I have elicited these respective intuitions using two hypothetical cases. From the *World Cup* case is the intuition that no number of claims weak claims can outweigh a much stronger claim. From the *Life versus Paraplegia* case is the intuition that some number of moderately strong claims can outweigh a very strong claim. One way for defenders of Pure Aggregation and Anti-Aggregation to regain the upper hand is to show that our intuitions in one of these cases are unreliable. The motivating cases are supposed to elicit morally important intuitions, but if these intuitions are elicited by imagining the cases incorrectly or by morally irrelevant or misleading features of the cases, then we have little reason to take them seriously.²¹

In §3.1 of this chapter, I defend the *World Cup* case intuition against various arguments that may be pressed by proponents of Pure Aggregation. In §3.2, I defend the *Life versus Paraplegia* case intuition against various arguments that may be pressed by proponents of Anti-Aggregation. Finally in §3.3, I will explore the possibility that the *World Cup* and *Life versus Paraplegia* intuitions are unlikely to be reconciled in the Relevance View because they rely on such different rationales.

3.1 The *World Cup* intuition

Here is Scanlon's *World Cup* case again:

Suppose that Jones has suffered an accident in the transmitter room of a television station. Electrical equipment has fallen on his arm, and we cannot rescue him without turning off the transmitter for fifteen minutes. A World Cup match is in progress, watched by many people, and it will not be over for an hour. Jones's injury will not get any worse if we wait, but his hand has been mashed and he is receiving extremely painful electrical

²¹ Because I take it that Pure Aggregation and Anti-Aggregation both gain support from the intuitions that they can elicit, I do not entertain general sceptical arguments about intuitions in ethics. I also set aside the particularist charge that intuitions in these cases do not generalise.

shocks. Should we rescue him now or wait until the match is over? (Scanlon 1998, p.235).

Pure Aggregation implies that there exists some number of fans for which it would be right to let the broadcast proceed at Jones's expense. However, many share the strong contrary intuition that rescuing Jones is the right thing to do no matter how many football fans would be frustrated if we were instead to stop the broadcast. The *World Cup* case, then, seems to show that weaker claims do not always have the potential to outweigh competing stronger claims. I now explore four arguments that seek to establish that this intuition in the *World Cup* case is unreliable.

3.1.1 What if the fans knew...

The first objection to the *World Cup* case intuition is that it is not eliciting intuitions to do with competing claims at all. When we imagine the *World Cup* case, we imagine that the fans would share our knowledge about Jones. If the fans knew about Jones, then they would hardly be able to enjoy the match knowing that it came at such a high cost to someone else. We are consequently imagining that everyone (including the fans) has an interest in stopping the broadcast and saving Jones. This means that we are not thinking of the *World Cup* case as involving competing claims at all.

This objection is question begging at best. Even if the fans did in fact know about Jones, their outrage could be because they recognise that their claims are not relevant to Jones's claim. Indeed, if the fans knew about Jones, and they did not show outrage, I would think that their enjoyment of the football match would be in a sense *inappropriate*. It would be inappropriate to value watching the football match when one grasps that doing so comes at such a high cost to another person. The Relevance View helps explain why it is inappropriate to value watching the game at the expense of another's suffering: what is at stake for the fans and what is at stake for Jones are poles apart.²²

3.1.2 Psychological biases

The Relevance View is not supposed to appeal to special duties that come from intimacy with the person with the strongest claim. Our duty to save Jones is supposed to come from just the difference between what is at stake for the fans and what is at stake for Jones. If our intuition to save Jones were

²² Cullity (2018) calls a related idea 'content undermining'. Cullity gives the example of sadistic pleasure not being the kind of thing we have reason to promote. I am not saying that the fans' pleasure in watching the game is in fact sadistic, but rather that if they knew about the situation then their pleasure would be tainted in a similar way that we think sadistic pleasure is tainted. This is different from the idea expressed by Voorhoeve (2014), discussed in the previous chapter, that any given fan's preference that Jones's claim be satisfied in a one-one competition is an apt test for relevance. As I noted, the threshold where one is morally required to prefer the satisfaction of another's claim may be lower than the threshold an agent would tolerate when adjudicating between competing claims. This means that Voorhoeve's test may act as a sufficient but not necessary condition for determining when one claim is not relevant to another.

being influenced by mere intimacy with Jones, then we would have good reason to discount the *World Cup* case intuition.

Hirose (2015, p.49-50) suggests that certain psychological factors may bias us in favour of satisfying Jones's claim. The first factor is that when we imagine the *World Cup* case, we might imagine being reasonably proximate to Jones. If we were instead in a room alongside a football fan on the other side of the world, we might feel less sympathy for Jones's claim. The second factor is the 'identified life' effect. For example, we are much more likely to give to Oxfam if we are told a story about a particular child than if we are presented with a relevantly related mere statistic. In the *World Cup* case, we are acquainted with Jones and his situation, but the football fans are presented as an anonymous mass. If we were instead to know the identities of the football fans, or we were to know only that *someone's* arm is trapped in electrical cabling, we might feel less empathy for Jones's claim and more empathy for the fans' claims.

I think many people's intuitions will remain unchanged even with these variations to control for our intimacy with Jones's predicament. Moreover, our empathy for Jones in these cases may be reliably grasping just how much Jones has at stake compared to any given fan.

Scanlon gives a more vivid illustration of a case that downplays our empathy for Jones:

Suppose, for example, that we are deciding whether to build a new system of transmitting towers that will improve the quality of reception for many television viewers. It may be highly probable that in the course of this project a number of workers will suffer harms at least as great as Jones's. Yet we do not think that it is therefore wrong to go ahead.
(Scanlon 1998, p.236)

This new case is analogous to the original case in that our decision to build the transmitting towers is justified by the minor benefits that it provides for a potentially great number of people but will come at the cost of one person suffering greatly. It is different from the original case in that we do not know who will suffer or even precisely when they will suffer. Even if you think that building these transmitting towers is still impermissible, this intuition will probably be less strong than the original intuition in the *World Cup* case. If this analogy has successfully neutralised our bias towards the person with the strongest claim, then perhaps this is a reason to discount our intuition in the *World Cup* case.

However, Scanlon notes that this case "involve[s] failing to prevent accidental injuries rather than [...] withholding aid from people who need it, in order to bring small benefits to others" (p.236). He also notes that "Our sense that it is permissible to undertake these projects also depends crucially on the assumption that precautions have been taken to make the work safe and that, in addition, workers have the choice of whether or not to undertake the risks involved" (p.236). I will say more about risk in the context of the Relevance View in Chapter 4 and 5. For now, I think this new case is sufficiently

different to the *World Cup* case and so it is not a more accurate illustration of the same case.²³ I conclude that psychological biases of the kinds we have been considering are not obviously problematic in the *World Cup* case.

3.1.3 Large number scepticism

One well-rehearsed sceptical worry is that our judgements in cases such as Scanlon's *World Cup* case are unreliable because we tend to have a poor imaginative ability to aggregate the claims of large numbers of stake-holders. Because we find it difficult to imagine the aggregate strength of millions of minor claims to watch the football match, we fail to conceive of our duty to leave Jones to suffer when the number of football fans is very large.

However, the judgement that one claim is not a relevant consideration is independent of our grasp of large numbers (Voorhoeve 2014, pp.75-6). The judgement that some claims are not relevant to others is not based on trying and failing to imagine a number of very weak claims that can outweigh a stronger claim. Rather, our judgement is that very weak claims are not the kinds of things that outweigh much stronger claims.

It is interesting to note that the sceptical worry being considered is also a popular move in population ethics. Total utilitarianism seems to entail the position that a large enough population where everyone has a life that is barely worth living is better than a smaller population where everyone has a life well worth living. This is known as the Repugnant Conclusion (Parfit 1984). Some claim that this conclusion would not, in fact, strike us as repugnant if only we could grasp the aggregate happiness of all these people whose lives are barely worth living. But again, it is not clear that our judgement (of "repugnance" rather than disrespect in this case) comes from our inability to properly imagine aggregate happiness of this hypothetical large population. There may be other explanations such as our concern that future generations flourish and the Repugnant Conclusion is a setback to this concern. We care deeply about people's quality of life and not just that they have a tolerable life. Total utilitarianism forces us to abandon this powerful intuition for the sake of theoretical simplicity.²⁴

²³ Here is another adaptation that might be thought to reveal that our intuition to save Jones in the *World Cup* case is problematically biased. Suppose that we knew that stopping the broadcast would lead to violent protests on the streets or lead one frustrated viewer to commit a murder. The consequentialist would perhaps view leaving Jones to suffer as the lesser of two evils. However, deontologists are well equipped to handle these kinds of problems. Deontologists are used to tolerating inefficiency by making use of agent-relativity. Agents have certain constraints on how they can act which are (at least partially) resilient to greater harms coming as side-effects of their actions. The fact that a murder will take place as a side-effect of helping Jones is a bad thing no doubt, but this *wrong* is not the wrong of the agent who chooses to save Jones. Whilst the deontologist should not tolerate too much inefficiency (e.g., 1,000 predicted murders might be too many), the *World Cup* case intuition is fairly resilient to this line of objection.

²⁴ See Scheffler (2018) for an exploration of this idea.

The *World Cup* case intuition and the Repugnant Conclusion do not show that our imagination is deficient, but rather they highlight the need to imagine richer possibilities underlying our moral thought. This need is met through the notion of relevance in the first case, and may be met by examining a richer set of attitudes towards future generations in the second.²⁵

3.1.4 Reliability strategy

The next objection I consider is that there is a certain practice that is reliable in the large majority of cases, but generates moral mistakes in ‘large number’ cases such as the *World Cup* case. For example:

[W]hen faced with the opportunity to save one person from being killed, it is often safe to ignore minor harms to others on the grounds that it is normally impossible to establish that enough people will be saved from minor harms to justify refraining from saving a person’s life given epistemic and time constraints on decision-making. (Tadros 2019, p.176)

This suggests that the judgement that certain considerations are not relevant could just be an artefact of efficient and generally safe moral decision making.

Tadros responds to this objection by arguing that the reliability of excluding minor harms from deliberation does not vary between ‘small number’ and ‘large number’ cases, because this practice is insensitive to such numbers. It is instead sensitive to a certain sort of disrespect involved in thinking that very weak claims are the kinds of things that can establish duties in contexts where they compete with strong claims. Go back to the analogy with consent. We treat consent as a decisive consideration not because considerations of autonomy usually outweigh considerations of welfare. Rather, violating consent shows a distinct kind of disrespect that is independent of how much a patient has to gain from being operated on.

This concludes the first major part of this chapter that investigates sceptical arguments against the *World Cup* case intuition. I have provided a response to each of these arguments. If, after a more comprehensive comparison between the merits of the Relevance View and Pure Aggregation, Pure Aggregation comes out on top, then we may have good reason to dismiss the *World Cup* case intuition. However, what I have shown is that the *World Cup* case intuition cannot be easily dismissed without a much wider-reaching analysis than I plan to entertain here.

3.2 The *Life versus Paraplegia* intuition

Of course, it is not only the proponents of Pure Aggregation that will resist the Relevance View; there is opposition from proponents of Anti-Aggregation as well. There are a number of instances where aggregation looks highly attractive. One extreme example is the choice between preventing one death,

²⁵ Draper (2002) has likened insistence in these cases on Pure Aggregation despite our intuitions to the contrary to a kind of “benefit fetishism”: showing more concern for maximising aggregate welfare than to the moral relationships between individuals.

and preventing 1,000 people from becoming paraplegic. It is highly intuitive here that we should serve the claims of the many individuals who each face a less significant harm than that faced by the worst placed individual. In this section I will investigate rationales for doubting the *Life versus Paraplegia* intuition.

3.2.1 The real impact of certain harms

The first argument I consider is that the *Life versus Paraplegia* case elicits the wrong intuition in a wide range of actual and possible cases of its type. One problem with the *Life versus Paraplegia* case is that it seems to assume that death and paraplegia affect all people the same way. In reality, paralysis may have a much larger impact on a champion runner than it would have on a person who works a desk job. For the champion runner, a major life-project is thwarted, and paraplegia may be as bad for them as death (or heaven forbid worse than death). Although the *Life versus Paraplegia* case invites us to imagine 1000 “very average” cases of paraplegia, it is difficult not to see this setting back people’s interests to vastly different degrees. When something as significant as paraplegia affects a large number of people, there is a reasonably high chance that this will affect someone as badly as death would. So, perhaps some claims against paraplegia are in fact as strong as the single claim against death.

However, even if the two competing groups of claims contain claims of the same strength, we still need some way of determining which group’s claims to satisfy. I find flipping a coin to be deeply counter-intuitive in this setting; saving the many surely has more pull than this. But perhaps as long as a few people affected as badly as death, all that is needed to support a duty to save the many is a tie-breaking argument. I expressed in the last chapter. I think tie-breaking is just a primitive kind of aggregation. It acknowledges that claims can cancel each other out and are in this sense commensurable with each other. Moreover, once we acknowledge claims are commensurable it is difficult to resist widening the scope of commensurable claims to claims that are only slightly worse than the strongest claim. So, even if we imagine that paraplegia affects some people as badly as death, the rationale for this judgement that we should help the many is still because the many outweigh the one.²⁶

²⁶ Wallace (2019, pp.216-9) makes a more promising argument for saving the many without appealing to aggregation. When competing claims are of the same strength, it is in the *ex ante* interest of everyone that the larger group to be saved because any given person is most likely to be in that group. By choosing to aid the few, the agent disrespects not only those in the larger group but also those in the smaller group who also had an *ex ante* interest in the larger group being saved. Wallace makes an helpful analogy: suppose I break my promise to drop you off at the airport and coincidentally prevent your death because your would-be plane crashed after take-off; even though you cannot justifiably regret my not picking you up, you have still been disrespected by me. Likewise, even though those in the smaller group will regret the agent’s decision to help the larger group, the agent still disrespects them by not aiding the larger group. Notice that this justification is only available when all claims are of equal strength because otherwise the *ex ante* justification would be overridden by the consideration to save the person with the strongest claim. While I find Wallace’s argument from *ex ante* rules when

Life versus Paraplegia is designed to be a *reductio* of Anti-Aggregation, but there is no denying that it is sufficiently under-described to let other assumptions influence our judgements. One such assumption is that certain individuals will respond very badly to paraplegia. However, rather than pursue this line of thought further, I will now explore a different strategy for scepticism about what the *Life versus Paraplegia* case shows.

3.2.2 Two domains of morality

The *Life versus Paraplegia* case is an extreme example that even staunch proponents of Anti-Aggregation should be moved by. However, there may be ways of accepting the intuition without accepting that any form of aggregation is the general upshot. After exploring examples that resemble *Life versus Paraplegia*, Wallace (2019) considers ways of making some concessions to interpersonal aggregation without giving up on his anti-aggregationist account.

First, Wallace notes that cases such as *Life versus Paraplegia* strike him as particularly “bureaucratic or administrative” in nature:

[T]hose bodies that are authorized to make such allocative decisions might naturally be understood to have special responsibilities, which they perhaps owe to the democratic populations on whose behalf they are acting, to see to it that public investments are made in ways that benefit as wide a population as possible. (Wallace 2019, p.228)

Our duty to save the many in the *Life versus Paraplegia* case may come from duties specific to bureaucratic bodies. Taurek (1977) makes a similar point when he suggests that where patients have some collective ownership of a resource, saving the many could be justified by agreeing in advance that the resource should be used to save as many people as possible. Both these suggestions mean that certain agencies are subject to further constraints placed on them as a public body that give them duties to bring about optimistic outcomes.

However, Wallace rightly worries that these aggregative considerations may not be so easily demarcated to another domain of practical reason. There may be *Life versus Paraplegia* style cases that do not occur within the domain of a bureaucratic decision with a democratic mandate. He also considers the idea that once a certain threshold of aggregate welfare is reached, this may operate as an *extra-moral* input into practical deliberation. Wallace contrasts extra-moral considerations with interpersonal moral considerations that he takes to be distinctive of the anti-aggregative approach. Extra-moral seems to be a kind impersonal moral consideration; something valuable but somehow divorced from the relations between persons (perhaps something like beauty).²⁷ When this pluralism of

competing claims are of the same strength interesting, I remain convinced that there are situations where competing claims are clearly not of the same strength and yet we ought to save the many who have weaker claims than the few.

²⁷ Johnson (2019) proposes that a threshold can be justified by acknowledging that individuals have two kinds of claims and so does not appeal to impersonal value. In the current context, the first kind of claim is held by the person with strongest claim whereas the second kind of claim is held by all

considerations—between anti-aggregative and aggregative considerations—is at play, Wallace suggests that these conflicts may not fully resolve themselves. So, even if we should and do serve the claims of the paraplegics, we have wronged the person who dies in some distinctly interpersonal way (Wallace 2019, pp.231-3).

Wallace’s two suggestions may allow him to hold onto the dominance of anti-aggregative reasoning by making a concession in extreme cases, but they also have the fault of compromising on parsimony: we must accept two domains of practical reason, one interpersonal and the other either “bureaucratic” or “extra-moral”. The Relevance View is committed to just one domain where claims compete. Moreover, the very idea of splitting morality into more than one domain in order to exclude certain cases looks troublingly *ad hoc*.

3.3 Hedging our bets

Building on Wallace’s idea from the last section, Pure Aggregation and Anti-Aggregation may just be two incompatible ways of looking at morality. Given moral uncertainty about which view is correct, we could play it safe by being guided by one or the other view in different situations (Temkin 2012, pp.39-45).

Even if we have good reason to believe that Pure Aggregation is correct, as long as we have some reason to believe Anti-Aggregation is also correct, there are certain circumstances where we should let Anti-Aggregation guide us. For example, in the *World Cup* case, so long as the total number of fans is not too large, our bets should plausibly be guided by the following considerations. If Anti-Aggregation is true and yet we choose to satisfy the claims of the fans in the *World Cup* case, we act *very* wrongly, because the claims in question are so far apart from each other in strength. However, if Pure Aggregation is true, and yet we choose to serve Jones’s claim, we do not act very wrongly, because the aggregate of the claims of the fans is not that much stronger than Jones’s claim. So choosing to save Jones would hedge our bets: we are right if Anti-Aggregation is true, and we are not far wrong if Pure Aggregation is true.

Where the verdicts of Pure Aggregation and Anti-Aggregation come sharply apart, we might have most overall reason to decide *as though* the Relevance View is correct. The Relevance View would be a way of hedging our bets: it minimises the chances of our making decisions that are very wrong.

While this pseudo-justification for the Relevance View might get proponents of Pure Aggregation and Anti-Aggregation to make some concessions, I think the justification of the Relevance View I outlined

parties. This weaker second kind of claim can outweigh the first only in very large numbers. One problem with this proposal is that it implies the number of claims is a peripheral issue for deontology which I argued was a mistake in §2.5.2.

in the previous chapter is powerful enough to resist this alternative justification.²⁸ The Relevance View should not be thought of as a hybrid between Pure Aggregation and Anti-Aggregation (*pace* Voorhoeve 2014) but rather seen as the limit of Weak Aggregation. As a burden gets more and more distributed, the aggregation of claims that each part of the burden grounds becomes weaker. As I noted in §2.7.2, Weak Aggregation embodies the value of solidarity: ‘we will each take one lashing to prevent our friend from receiving ten’.

The responses I have outlined in this section and the last are worthy of much more development than I can give them here. However, my aims are more modest than showing that these alternatives are false all things considered. I am trying to make the best case for the Relevance View, and the aim of this chapter is to show that the intuitions that motivate the view are more difficult to resist than some might think.

3.4 Conclusion

In this chapter, I have explored a number of rationales for doubting the intuitions that motivate the Relevance View. In each case I have offered a response in defence of these intuitions. A more worrying challenge to the Relevance View is that by helping itself to both Pure Aggregation and Anti-Aggregation, it has the theoretical support of neither. I explored the idea that Pure Aggregation and Anti-Aggregation are incompatible ways of resolving competing claims. However, this thesis is centrally concerned with making the best case for the Relevance View rather than fully responding to this irreconcilability.

The Relevance View may be able to capture more of our intuitions than either Pure Aggregation or Anti-Aggregation can do alone, but by forging a middle way between these two extremes it faces a unique set of problems. The next two chapters explore how the Relevance View might behave in various real-world settings. The problems explored in Chapters 4 and Chapter 5 threaten to show that practical adherence to the Relevance View in different settings would make everyone worse off in counter-intuitive ways. Chapter 6 explores how the Relevance View can respond to choices that involve more than two homogenous groups competing for aid. Whereas the problems explored in Chapters 4 and 5 are solely inherited from the upshot of *World Cup* case type intuitions, the problems in Chapter 6 come from trying to reconcile the *World Cup* and *Life versus Paraplegia* intuitions.

²⁸ We could, of course, endorse both rationales. But this depends on whether or not one thinks there is one true justification for a theory, or theories are always justified by weaker, perhaps more pragmatic, considerations.

Chapter 4

Intrapersonal Aggregation and the Unity of Individuals

In the previous chapter, I identified several sceptical arguments that could be levelled against the Relevance View on behalf of proponents of Pure Aggregation and proponents of Anti-Aggregation. In each case, I offered a reason to think our intuitions in the *World Cup* and the *Life versus Paraplegia* cases are reasonably reliable. However, one further potentially devastating line of argument against the Relevance View comes from the fact that acting in accordance with it would run counter to many of our practices. This chapter begins the second general task of this thesis which looks at how the Relevance View might behave in real-world settings. My general strategy is to restrict the scope of the Relevance View to avoid counter-intuitive results. For each such restriction, I will offer a rationale that it is not *ad hoc*. In this chapter, I make the first step towards minimising how revisionary the Relevance View would be by excluding its application to cases where acting in accordance with it would make everyone worse off in counter-intuitive ways. I justify this restriction on the grounds of the Unity of Individuals, the counterpart to the Separateness of Persons.

The challenges in the previous chapter arose because the Relevance View resembles Pure Aggregation in some contexts and Anti-Aggregation in other contexts. The challenges in this chapter and the next are come from the part of the Relevance View that resembles Anti-Aggregation. These problems seem to arise from not giving equal marginal weight to each claim of a given strength. Therefore, taken together, many of these arguments may be seen as an additional reason to doubt the *World Cup* case intuition.

In §4.1 I will show how intrapersonal aggregation interacts with a number of distributive justice decisions. In §4.2 I will argue that the Unity of Individuals explains why Relevance View would not apply when acting in accordance with it would make everyone worse off. In §4.3 and §4.4 I engage in an extended discussion of the Relevance View and risk. The Relevance View must tread a fine line between tolerating risk on the grounds of intrapersonal considerations and not tolerating so much risk as to undermine itself. These sections will also establish the assumptions for the discussion in Chapter 5.

4.1 Intrapersonal aggregation

Take the following sample of policy decisions:

- A government agency's decision to fund one kind of treatment or drug over another.
- A government agency's decision to build a highway.
- A Public health intervention such as a sugar tax.

Decisions of these kinds inevitably involve trade-offs between the interests of different people. More than that, these decisions often involve trading weak interests against much stronger interests. For example, a government agency may have a choice between funding paracetamol for mild pain relief and funding a lifesaving treatment for some rare disease. If paracetamol is much cheaper than a lifesaving treatment, then we can expect to benefit many more people by funding the former. However, claims to mild pain relief may not be relevant to claims to lifesaving treatment, and funding paracetamol here may violate the Relevance View.

Another example of a public policy decision that seems to violate the Relevance View involves building a highway that will provide a benefit to many people, but will put the residents near the highway at higher risk of early death. For example, children from nearby may walk out onto the road and all local residents are victims to the detrimental effects of exhaust pollution on their health. The Relevance View might condemn building this highway if it involves genuine competing claims and the minor benefits that drivers can expect are not relevant to these major harms that may befall local residents.

One consideration that seems to matter in both of these examples is whether or not the people whose claims are thwarted by acting one way are the same people whose claims are satisfied. The *World Cup* and the *Life versus Paraplegia* cases are set up as distinctly interpersonal trade-offs. If we satisfy the claims of one set of individuals, we frustrate the claims of another set of individuals. *Interpersonal* trade-offs can be contrasted with *intrapersonal* trade-offs. Intrapersonal trade-offs involve trading harms and benefits within a given individual's life. Because individual lives exhibit a morally important unity, we do not usually view trade-offs within lives as involving different claimants. My decision to have a hip replacement now will thwart my short-term interests, for example, but will increase my quality of life in the coming years. It would be odd to view this as a trade-off between my current self and my future self, as akin to a trade-off between two different individuals. This odd approach is the error that is identified by the Separateness of Persons.

Where an act or series of acts can make every individual better off than an alternative act or series of acts, this provides an independent justification for acting in this way. Moreover, as I will argue, this independent justification may be overriding.

4.2 Iterated choices

One reason that the treatment of minor ailments in a public healthcare system may be consistent with adherence to the Relevance View is to do with intrapersonal aggregation. If a choice occurs frequently enough, then it may be in everyone's long-term interest to restrict the scope of the Relevance View.

The following table is a simple illustration of how being guided by the Relevance View at each point in time can make everyone worse off than acting in a way that maximises every individual's own welfare. In this case, there are four individuals, *P1-P4*, and four times where claims compete, *t1-t4*.

Suppose claims of strength two are not relevant to claims of strength five, and that at each time an agent can either satisfy the multiple claims of strength two or the single claim of strength five, but not both. For instance, at $t1$, the claims of $P1$ - $P3$ are each of strength two and compete with the claim of $P4$ that is of strength five. If we apply the Relevance View to this decision, then we should satisfy $P4$'s claim.

Table 1 An Iterated Choice

	$P1$	$P2$	$P3$	$P4$
$t1$	2	2	2	5
$t2$	2	2	5	2
$t3$	2	5	2	2
$t4$	5	2	2	2
<i>Total when guided by Relevance View at each time</i>	5	5	5	5
<i>Total when guided by considerations of intrapersonal aggregation</i>	6	6	6	6

By applying the Relevance View at each time, everyone has a claim of strength five satisfied. But this comes at the cost of three other times when a claim of strength two is thwarted. The aggregation of these thwarted claims is of strength six, which is greater than acting on the basis of the Relevance View at each time.

Intrapersonal aggregation provides a justification for not acting in accordance with the Relevance View at each time.²⁹ This means that the choice situations where the Relevance View does apply must be rare enough such that intrapersonal considerations do not dominate. This suggests the following restriction on the scope of the Relevance View:

An agent should not be guided by the Relevance View when acting in accordance with it would make *everyone* worse off.

The rationale for this restriction comes from the counterpart to the Separateness of Persons. People's lives have a unity that makes intrapersonal trade-offs appropriate in a way that interpersonal trade-offs may not be appropriate. The Unity of Persons shows why the Relevance View would not apply in situations where acting in accordance with it would make everyone worse off in the long run.³⁰

4.3 *Ex ante* reasoning

There is a more liberal sense in which strict adherence to the Relevance View seems to make everyone worse off: the Relevance View may run counter to everyone's *ex ante* interests. *Ex ante* reasoning takes the temporally *prior* perspective. This can be illustrated using *Table 1* again. Prior to any time at which claims compete, it is in each of $P1$ - $P4$'s *ex ante* interest that the agent satisfies the set of

²⁹ See Temkin (2012) for a discussion of how intrapersonal aggregation faces similar obstacles.

³⁰ Because everyone would want the Relevance View to give way to pareto-optimality when these two views give conflicting verdicts, we could even assume that everyone would waive their claim at each particular instance if it meant that they would be better off in the long run.

weaker individual claims. Before we know who among $P1-P4$ will have which claim at $t1$, each person has a 1 in 4 chance of having a strength five claim ($0.25*5 = 1.25$) and a 3 in 4 chance of having a strength two claim ($0.75*2 = 1.5$). This means that each person has the most interest in the agent acting against the recommendation of the Relevance View at $t1$. Identical reasoning supports acting in everyone's *ex ante* interest at each of $t2-t4$.

4.3.1 Productive social risk

Ex ante reasoning looks particularly attractive in certain settings. Take the following hypothetical case:

One million children have contracted a virus. If the public health agency does nothing, all one million children will be ill for one week and then die. The public health agency can produce one of two vaccines that must be taken by all one million children to be effective.

Vaccine 1 is certain to save the life of each child but will do nothing for their being ill for one week.

Vaccine 2 will prevent each child from being ill for one week and save their life, apart from a 1 in 100,000 chance that it does nothing.³¹

In this case, doing nothing is morally out of the question. Any plausible theory of distributive justice would agree with that. I would also argue that by modest assumptions about self-interest, it is in any given child's self-interest to receive *Vaccine 2* rather than *Vaccine 1*. Notice that we take risks of this nature all the time. When we take a walk to the shop, we incur a small risk of significant harm (such as being run down by a car) for the sake of a high chance of a relatively insignificant benefit (such as getting a chocolate bar). Living with extreme risk aversion is a sure recipe for living an impoverished life.

If it is reasonable that each child would choose *Vaccine 2* for themselves, then it is highly plausible that we as the agent should choose *Vaccine 2* for them all. However, if one million children receive *Vaccine 2*, we can be almost certain that one child will die.³² A claim not to be ill for one week is plausibly much weaker than a claim against death. If claims against death are competing against claims against illness, then according to the Relevance View, a week of illness is not relevant to death,

³¹ Adapted from Dougherty (2013) and Frick (2015). Frick (2015) defines Social Risk according to four criteria: "(1) The risky action or omission will affect a large number of individuals. Because of this, it is virtually certain that some people will end up being burdened by it. (2) The individual losses to those who are burdened (relative to the baseline of some available alternative) are considerably greater than the individual gains for those who are benefited. (3) The action-type in question is rare, or rarely affects the same people twice; as a result, we cannot assume that over time almost everyone will benefit from a principle that permits actions of this type to be performed. (4) The risky action or omission is intuitively permissible" (p.178-9).

³² The chance that at least one child will die is $1 - (1/100,000)^{1,000,000} \approx 1$. Matters are slightly complicated when we are not privy to a precise risk. See Rowe & Voorhoeve (2018) for a discussion of decision making under 'severe uncertainty'.

and so we ought to satisfy the claim against death. This would require us to choose *Vaccine 1* where everyone is ill but nobody dies.

Of course, before an agent acts, she does not know who will die, but she can be quite certain that *someone* will die. Nonetheless, she can still deliberate on the basis of what she has good reason to believe will eventuate. This is to take the *ex post* perspective. The *ex post* perspective can be contrasted with the *ex ante* perspective. From the *ex ante* perspective, each child has an interest in *Vaccine 2* that is full health discounted by a 1 in 100,000 chance of death.

One way of respecting the Unity of Individuals is to take productive risks on their behalf as they would take for themselves. The *ex ante* perspective is the perspective each child takes when they choose *Vaccine 2* for themselves, it also rationalises our choice to administer *Vaccine 2* on their behalf. Evaluating moral choices on the basis of *ex post* claims would threaten the symmetry between the prudence associated with individuals taking risks and beneficence associated with taking risks on behalf of others. Taking the *ex post* perspective would instead mean that we should choose for all children differently from how we would choose for them each individually (Frick 2015).

One of the reasons we should favour the *ex ante* perspective, in this particular case, has to do with intra-personal aggregation again. We can use a similar justification to limit the scope of the Relevance View as we did when it made everyone worse off. Instead of trading-off between claims of different individuals, we are here trading-off within the possible futures of the same individuals (Thompson 2018). Part of respecting our Unity as Individuals is letting people run risks and taking certain risks on the part of others. Because the Relevance View is only concerned with inter-personal aggregation, it is not appropriate to apply the Relevance View in settings where taking a risk is in the interests of everyone.

4.3.2 Does *ex ante* reasoning undermine the Relevance View?

One issue with the justification offered at the end of the previous section is that it actually seems to undermine the Relevance View. We could reason that it is often in everyone's *ex ante* interests to follow policies that maximise aggregate welfare rather than comply with the Relevance View. This is because we often do not know who will be unlucky enough to have very strong claims to aid in the future. Suppose that letting the broadcast continue in the *World Cup* case would lead to more aggregate welfare than saving Jones would. We could justify letting Jones suffer in the *World Cup* case by pointing out that this policy was in his own long-term self-interest (his self-interest behind the veil of ignorance, to use a popular image). It may be in our *ex ante* interest to risk an infinitesimal chance of finding oneself in Jones's position rather than an almost certain chance of having one's favourite sports match disrupted.

However, the foregoing would completely undermine the Relevance View and gets embarrassingly close to certain developments of utilitarianism.³³ This kind of *ex ante* reasoning would also be devastating for deontology in general. To use another example, it may be in everyone's *ex ante* interest to endorse a society that adopted a policy of killing an individual to harvest their organs where those organs could save more than one life. This would, after all, decrease everyone's expected chances of an early death. However, such reasoning would completely undermine the deontological distinction between killing and letting die. In order to save the Relevance View I need to differentiate productive social risk from problematic *ex ante* rules.

In order to restrict the reach of *ex ante* justifications, Frick (2015) argues that their use must pass the following test:

Decomposition Test: If a rule or procedure can be decomposed into a sequence of distinct causal stages, each of which involves the voluntary action of some agent, then it is permissible to adopt and act on this rule or procedure only if the actions it requires at every stage are justifiable to each person *at that time*. (Frick 2015, p.205)

Frick is writing in the context of contractualism, but his point generalises. The core idea is that *ex ante* justifications must be present at every stage of a procedure in order for it to be permissible. We cannot justify leaving Jones to suffer at the time of the incident merely because this rule maximises his long-term *ex ante* interests. (Perhaps it is in Jones's power to waive his claim to be saved, but this does not give him a duty to waive his claim.) Jones has a valid claim to being saved in the *World Cup* case at the time the case is set. What allows us to administer *Vaccine 2* to the population of children is that it is in each of their *ex ante* interests at the time of action. There is no later point in time where individuals will have claims on us that we act otherwise. This is because by then it will be too late. After giving all children *Vaccine 2* and identifying the children who will die, it is impossible to reverse this decision and give them all *Vaccine 1*.³⁴

The Relevance View protects us against having our comparatively strong claims defeated by the much weaker claims of others. This protection is the default. To assume otherwise would concede too much to the utilitarian. All that is needed to satisfactorily defend agents' duties to take risks on behalf of patients without undermining the Relevance View is to say in certain circumstances that patients are not protected by this default: when an agent can act in everyone's *ex ante* interest at each time, then the agent has a permission to act in this way.

³³ For example, the rule utilitarianism of Harsanyi discussed in Frick (2015, p.189)

³⁴ However, if it were possible to switch the vaccine at this time, then this choice may be morally different. Gordon-Solmon (2019) has recently disputed Frick's decomposition test on the grounds that it cannot differentiate between decomposable acts and acts that look identical to these that are set up to make it impossible to decompose them. I recognise more work needs to be done in response to Gordon-Solmon's work but developing such a response lies outside the scope of this thesis.

Walen (forthcoming, p.6) points out that “if we are to have our own lives to lead, we have to run our own risks; we cannot have our welfare automatically pooled for greater *ex ante* benefit”. Having one’s organs taken for the benefit of another person is a huge sacrifice and so unless we have agreed to this in advance, morality that respects our autonomy should not expect us to be committed to *ex ante* rules (Walen 2019, p.62).

There is another Separateness of Persons justification for the move made in the previous paragraph. Our bad luck falls on us as individuals, not as some supra-individual deliberating behind the veil of ignorance. The decomposition test thus protects our unity as individuals by allowing others to take risks on our behalf without undermining the Relevance View. We can sometimes act in someone’s *ex ante* interest in the same way that they themselves would choose to act under normal circumstances, but this does not commit us to any far wider deployment of *ex ante* justification such as utilitarianism.

(See Appendix A on an alternative strategy for defending the Relevance View without undermining individual autonomy. This strategy relies on differentiating between acts and conjunctions of acts).

4.4 Two further problems with restricted *ex ante* reasoning

I will now consider two further problems that will help refine the scope of the Relevance View in relation to *ex ante* reasoning.

4.4.1 The problem of dependent risk

In the social risk case we have been considering, everyone has an independent risk of having a fatal reaction to *Vaccine 2*. Otsuka (2015) gives a helpful analogy to the distinction between dependent and independent risk by asking us to imagine two different roulette wheel setups. First, imagine a giant roulette wheel that has one million slots over a city. Whichever slot the ball lands on will correspond to death of one particular person. In this case, the risk is dependent: my not getting hit by the roulette ball *depends* on someone else’s being hit by that ball. Call this case ‘whole-city roulette’. Compare this with a case in which each city dweller has their own roulette wheel over their head with a million slots and one ball. This time, even though the expected number of deaths is the same as in the previous case, my not getting hit by my ball has nothing to do with whether or not you get hit by your ball. Call this case ‘individual roulette’.

In the vaccine case we have been considering, everyone has an independent risk of having a fatal reaction to the vaccine. All we know is that one’s susceptibility to this is 1 in 100,000; taking *Vaccine 2* is more like ‘individual roulette’ than ‘whole-city roulette’. However, consider a new set-up where *Vaccine 2* is replaced with *Vaccine 3*:

One million children have contracted a virus. If the public health agency does nothing, all one million children will be ill for one week and then die. The public health agency can produce one of two vaccines that must be taken by all one million children to be effective.

Vaccine 1 is certain to save the life of each child but will do nothing for their being ill for one week.

Vaccine 3 is certain to save the life of each child and also to prevent them from being ill for one week.

In this case we should choose *Vaccine 3*. It is identical to *Vaccine 1*, except that it also cures each child's illness. However, suppose just before *Vaccine 3* is administered, we realise there has been a dosing error: only 999,990 effective vaccines have been created, ten vaccines short. We can either insert 10 mock-vaccines at random into the batch and go ahead with administering *Vaccine 3* or switch to administering *Vaccine 1*.³⁵

By administering *Vaccine 3*, everyone still has a 1 in 100,000 chance of dying, which is identical to the chance of dying if one receives *Vaccine 2*. However, this time, unlike the original case, the risks are dependent. That is to say that the likelihood of any particular child's receiving the fully effective vaccine now depends on others receiving the fatal mock-vaccine. This was not the case in the original set-up: whether or not any particular child has a fatal reaction to *Vaccine 2* does not depend on other people having this reaction. We may seek to put ourselves at ease by noting that risks of this nature are unlikely to occur. After all, I have taken quite some creative liberty in designing this case. But for the sake of theoretical completeness, I will try to say something more.

This new case seems like the paradigmatic case of competing claims. Even though any single person's risk of dying from the vaccine is the same as the previous case, if one child benefits when the vaccine is rolled out, their benefiting from the vaccine comes at the cost of another being harmed.

(Alternatively, the *chances* of my being harmed are increased because somebody else is benefited).

Even though from the *ex ante* perspective, *Vaccine 2* and *Vaccine 3* are identical, there seems something uncomfortably interpersonal present in *Vaccine 3* that is not present in *Vaccine 2*.

One potential difference between the dependent risk and independent risk cases is that when risk is independent, there is a chance that nobody will die and so we may wish to hedge our bets on this possibility actualising. However, this is a dead-end. Not only is it incredibly unlikely that nobody will die in the independent case, but there is also the chance that even more people will die in the independent case than in the dependent case. After all, over two coin-tosses, the chance of flipping no tails is the same as the chance of flipping two tails.

Rather than say that the Relevance View should differentiate between *Vaccine 2* and *Vaccine 3*, the difference between these vaccines may involve an independent consideration about fairness. Rowe (2019) thinks that cases of independent risk are importantly different from cases where one benefits at another's expense.³⁶ It is only when everybody's *ex post* interests are independent that they do not

³⁵ See Ruger (2018) for a different set up that elicits a similar problem for contractualism.

³⁶ Rowe (2019, p.55-6 fn.20) admits that consent looks even more important in cases of dependent risk than in cases of independent risk. This might be worth exploring. On another note, Tadros (2019)

compete. One reason Rowe (2019, p.57) thinks that dependent risks are morally distinct is that “[g]aining at the expense of others is at odds with the value of solidarity”, which is “the desire for all to get something.” Whether the agent administers *Vaccine 1* or *Vaccine 3*, the children’s “fate is tied together.” If *Vaccine 1* is selected, then all children end up ill for one week. If *Vaccine 3* is selected, there is an inevitable inequality of outcome. Contrast this with *Vaccine 2* from the original case. In *Vaccine 2*, the inequality in outcome is resolved by the fact that everyone’s risk was independent; one child’s fate has nothing to do with the fate of the other children.

The feature of dependent risk may make *Vaccine 2* preferable to *Vaccine 3*, but this difference is not enough to make selecting *Vaccine 3* impermissible all things considered. I think this for the following reason. Suppose that everyone will receive *Vaccine 1* by default, but each individual can opt out and receive *Vaccine 3* instead. It is reasonable for every individual alone to opt out in favour of *Vaccine 3*. By opting out of *Vaccine 1*, any given individual does not affect the prospects of any other individual, nor what it is reasonable for any other individual to choose.

4.4.2 The problem of the identified life

I will now consider a second worry with using the *ex ante* perspective in cases of involving risk. After defending the *ex ante* perspective in the context of Scanlon-style contractualism which is a kind of Anti-Aggregation, Frick (2015) identifies a mixed blessing from adopting this perspective. Consider the following case:

Miners (1 vs. 100): Gareth, a miner, is trapped in a collapsed shaft. If we do not save him, he is virtually certain to die within days. However, a rescue will be costly. Suppose we must choose between the following two options:

- Rescue: Spend all our available funds to rescue Gareth.
- Prevention: Spend our available funds to improve safety at this mine, reducing the risk of future accidents. If we choose this option, the risk of death for each of the other 100 people working at this mine of dying in a future accident will be reduced from 3 percent to 1 percent. We expect that this will save two lives (though we cannot know whose). However, Gareth will die. (Frick 2015, p.215)

Frick believes that the *ex ante* perspective in combination with Anti-Aggregation gives us decisive reason to rescue Gareth in this case. This is because whereas Gareth has a full claim against death, each of the other miners have weaker claims that is discounted to reflect the 2 percent decrease in this death eventuating. The Relevance View is not a form of Anti-Aggregation, but when one claim is not relevant to another, it behaves like an Anti-Aggregationist view. Suppose a full claim against death is not relevant to a 2 percent increase in the chance of death. Then it seems that we should satisfy

writes about the importance of relevant alternatives. To illustrate this idea using the current cases: when *Vaccine 1*, *2*, and *3* are all options, *Vaccine 3* might be impermissible but not when just *Vaccine 1* and *3* are the only alternatives.

Gareth's claim no matter how many miners are thereby disadvantaged by having to deal with a high chance of death (3% rather than 1%).

The *ex ante* perspective in combination with the Relevance View seems to support favouring what Frick calls the 'identified life' over any number of 'statistical lives'. On the one hand, some have thought that preference for the identified life is a feature of common-sense morality and this would seem to count in favour of the view. On the other hand, when the number of miners is increased, using *ex ante* reasoning becomes less plausible. Suppose that the alternative to saving Gareth is now to reduce the chances of death for 1000 miners from 3 percent to 1 percent. We could now expect to save 20 statistical lives! Now saving Gareth rather than improving safety for 1000 miners should be much harder to justify. But *ex ante* reasoning in this case seems insensitive to the number of future deaths. This should strike the supporter of aggregation as deeply counter-intuitive.

Frick argues that the solution to the problem of identified over statistical lives is making some further concessions to Pure Aggregation. However, this significantly compromises the aims of his own Anti-Aggregation view and would do the same for the Relevance View. Instead I believe this problem can be solved by getting clear about the relationship between the Relevance View and *ex ante* reasoning.

The *ex ante* justification only works when everyone's claims are aligned. When claims compete, it makes more sense to use the *ex post* perspective (Walen forthcoming).³⁷ To justify the transition from the *ex ante* perspective to the *ex post* perspective consider the concept of waiving one's claim. There is an analogy between waiving your claim in a single-person case and waiving your claim in a case of social risk. If a patient consents to a surgery and the doctor takes reasonable precaution, the patient has no grounds for complaint if the surgery goes bad. By analogy, when we act in the *ex ante* interests of all one million children by giving them each *Vaccine 2*, it is reasonable to assume that each child will waive their *ex post* claim (Walen forthcoming). If a child dies, they have no complaint because we acted in each child's *ex ante* interest. However, if a course of action is not in someone's *ex ante* interest, we can hardly assume that they have waived their *ex post* claim. In the mining example, why would any of the 1000 miners waive their *ex post* claim against death just because another person (Gareth) faces a certainty of death? This might be a noble thing to do, but not something that can be reasonably assumed. The miners would only waive their *ex post* claim if we were acting in their *ex ante* interest.

Someone could object that the analogy to the surgery case and the vaccine case is weak. The *ex post* claims of the children who die are problematic precisely because they are caught up in an interpersonal trade-off with the claims of all the other children who will merely suffer illness for a week. By contrast, there is nothing interpersonal about the surgery case; the patient's interests alone are at stake.

³⁷ Several writers have developed hybrid risk models to make the Relevance View work across all cases involving risk (e.g., Lazar 2018; Suikkanen 2019). I agree with Walen that this is a general mistake in the literature on risk, but I will not engage with this line of argument here.

Perhaps a better analogy then would be the risk we impose on each other by driving our cars. The state has an interest in stopping individuals harming each other (this is a relational claim we have against each other). But we think that it would be wrong for the state to ban driving just because some people will seriously harm others. We all waive our complaints against being injured by other drivers so that we can all enjoy the benefits of driving (of course we waive these complaints assuming that due care is shown by other drivers).

I conclude that the *ex ante* justification does not support preference for the identified life. When *ex ante* interests compete we should revert to the *ex post* perspective for aggregating competing claims.

4.5 Conclusion

In this chapter, I have identified two restrictions on the scope of the Relevance View that come from considerations about intra-personal aggregation. The Relevance View can respect the Unity of Individuals by not applying in situations where it would make everyone worse off. Intrapersonal considerations also justify productive social risk. Another reason that this is important is to respect people's freedom to take individual risks. However, this does not justify the general adoption of *ex ante* rules. We can restrict the scope of application of *ex ante* reasoning using something like Frick's decomposition test, according to which *ex ante* justifications must be available at every stage in the series of acts in question.

In Appendix A I consider an alternative strategy for responding to worries about risk in relation to the Relevance View. I argue that this strategy is ultimately unsuccessful. In Appendix B I also consider somewhat tangential epistemic worries about *ex ante* justifications in situations involving risk.

Still, even with the arguments I have made in this chapter that appeal to intra-personal aggregation, the instances in which intrapersonal aggregation kick in remain very rare. The chances of everyone in the population having a major claim to aid at least once in their life are low. Nor does productive social risk exhaust the instances where the Relevance View seems to rule that many of our practices are impermissible. In the next chapter, I aim to refine the scope of the Relevance View even more by acknowledging two additional factors that may be relevant in real-world settings.

Chapter 5

How Revisionary is the Relevance View?

In the previous chapter, I argued for two domains that lie outside the scope of the Relevance View's application. First, the Relevance View would not apply when it would make everyone worse off in the long run. Second, the Relevance View would not apply in cases where taking a social risk is in everyone's *ex ante* interest at all times. Both of these cases are more appropriately understood as involving intrapersonal rather than interpersonal trade-offs. This distinction is grounded in the Unity of Individuals, the counterpart to the Separateness of Persons. Respecting the Unity of Individuals means prioritising trade-offs that make individual's lives as a whole go better. This includes taking prudential gambles on their behalf; after all, excessive risk-aversion is a sure way to lead an impoverished life. However, even with these limitations to the scope of the Relevance View, there are still some settings where the Relevance View looks excessively revisionary. Of course, the mere fact that acting in accordance with the Relevance View runs counter to our practices might just mean that our practices are wrong. However, this might also be a *reductio* against the view, so in this chapter I aim to show some ways that supporters of the Relevance View can sleep easy with many of our practices.

I first look at what the Relevance View would mean for setting priorities in healthcare settings. On the face of it, the Relevance View seems to endorse moving resources from the treatment of minor health burdens to major health burdens. I will look at some possible reasons why this may not be the case. Next I will look at what the Relevance View would mean for practices where one group imposes a risk on another group with the potential to cause individuals in that group significant burdens; the Relevance View seems to rule out such practices. I will argue that deontology has the resources to show why this is not the case.

5.1 Priorities in healthcare settings

Publicly funded health agencies provide health-related resources with limited budgets, so at some level, they must be engaging with competing claims situations. For example, most publicly funded healthcare agencies devote funding to the treatment of minor ailments such as broken fingers, as well as to more significant conditions such as heart failure. In theory, the funding for minor ailments could be redirected to the funding for major health burdens. Supposing that claims to have broken fingers fixed are not relevant to having one's serious heart condition treated, the Relevance View would seem to require redirecting these resources. A related example is when government agencies purchase drugs with a limited budget. A decision to choose to fund antihistamines for hay fever, for example, when these funds could go towards expensive drugs for rare life-threatening diseases, looks as though it might violate the Relevance View.

In these examples, acting in accordance with the Relevance View seems to require huge restructuring of how we manage healthcare.³⁸ Of course, there may be diminishing marginal returns by redirecting this funding; we will be able to increase aggregate well-being less and less as resources are transferred. However, the Relevance View is built on a rejection of total focus on efficiency when it comes to resolving competing claims.

Would acting on the basis of the Relevance View mean that we have to stop fixing broken fingers and funding antihistamines? I think an affirmative answer could be a *reductio* against the Relevance View. Surely the fact that minor burdens are so numerous and often require minimal resources to ameliorate, justifies resourcing their treatment. I now explore ways that defenders of the Relevance View could avoid these implications. In §5.1.1 I will outline and appraise resources that I have already mentioned in the preceding chapters. In §5.1.2 I will outline and appraise new resources that defenders of the Relevance View could bring to bear on this worry.

5.1.1 Defending the treatment of minor ailments: aforementioned resources

I have already said two things that may limit the reach of the Relevance View in the settings outlined above. First, in §3.1.2, I considered an example where we are faced with the decision of whether or not to build a transmitting station where we know that a construction worker is likely to be injured at least as severely as Jones in the *World Cup* case. In response to this adaptation of the *World Cup* case, I mentioned that we should distinguish between failing to prevent foreseeable harms and withholding aid that would ameliorate these harms. If we fund a new anti-histamine before spring and in spring a few patients present with major health burdens that could be ameliorated with another drug, we may now have no means of purchasing this second drug and so this looks like a situation where we are merely failing to prevent foreseeable harms. Perhaps there is a morally important difference between failing to fully resource a certain treatment and refusing to offer this treatment.³⁹

The second factor that may limit the revisionary character of the Relevance View in public healthcare settings is intra-personal considerations of the sort noted in the previous chapter. The Relevance View plausibly does not apply when acting in accordance with it will make everyone worse off in the long run. However, while treatment for health-related burdens is common, it is not so common that we can expect everyone's long-term interests to be identical. Some individuals are just genuinely less lucky than others, and will be burdened in ways that other people will avoid. Therefore, I do not think that this factor will cover all of our bases.

³⁸ Of course, perhaps the Relevance View rationalises the fact that actual practices of our health system do actually embody a lot of prioritisation of attention to major health burdens over attention to minor ones. I will return to this idea in §5.1.2.

³⁹ I worry that this response endorses a kind of practical irrationality, where we can make it impossible for ourselves to act later in the right way. Still, I will not explore this line of thought further here.

Another complication with relying on intrapersonal considerations is that an agent making resource prioritisation and allocation decisions will often not *know* whether or not the adoption of a policy will make everyone worse off in the long run. On the one hand, an agent may act in accordance with the Relevance View at $t1$ not anticipating that a similar situation in that domain will occur at time $t2$, and so on. On the other hand, the agent may anticipate that the choice set-up at $t1$ will occur frequently enough such that they can act in accordance in everyone's intra-personal interests, and what in fact happens is that the set-up at $t1$ never occurs again. With this uncertainty, the agent can be guided by what she has most reason to believe. If the agent has good reason to adopt a policy of responses and this policy turns out to be incorrect, then she still acts correctly in one sense.

I have mixed feelings about how useful the two factors I have just mentioned will be for justifying policies for the treatment of minor ailments. In light of this, I will now explore two reasons to limit the scope of the Relevance View that are independent of the considerations just outlined.

5.1.2 Defending the treatment of minor ailments: new resources

Voorhoeve (2018) suggests that the Relevance View is less revisionary once we consider that minor ailments might be the first stages of significant harm. For example, "a cold, if untreated, can develop into a life-threatening lung infection" (p.154). One's claim to have one's cold treated would then need to reflect this much larger setback to health-related well-being. Lung infections, unlike mere common colds, may now be relevant to more significant harms. The same could be said of broken fingers. If left untreated, a broken finger may impact on one's ability to engage in certain occupational and recreational activities in the future. This is not a matter of intrapersonal aggregation because we are not adding up potentially different claims over the period of one's lifetime. Rather, this factor merely recognises that the true strength of a claim reflects the total impact of well-being across one's lifetime and not merely the most immediate effects on one's well-being.

A second justification for treating minor ailments is that the treatment of minor ailments may not compete with the treatment of larger harms if the former are mandated on different grounds to the latter. Suppose, for example, that the chances of any given person holding a very significant claim to aid is rare enough, and the chances of any given person holding a minor claim to aid is common enough. Then individuals may prefer to gamble on not receiving aid when they have a significant claim, so that their minor claims are taken due care of.

The preference outlined in the previous paragraph may be expressed in multiple ways. For example, J. Paul Kelleher (2014) suggests that autonomous preference is an independent consideration that publicly funded health agencies can be sensitive to:

[If a government health agency] can make a good case that it is displaying enough benevolent concern for the lives and well-being of those it serves, it can then permissibly add a layer of policy aimed at satisfying individuals' autonomous preferences for risk-taking. (p.406)

By “benevolent concern”, Kelleher means acting in such a way as to satisfy the claims of those facing major harms. One problem with this solution is that “displaying enough benevolent concern” seems likely to be arbitrary. However, this might be more or less a matter of public deliberation.⁴⁰ How much do we as a society want to give priority to satisfying the strongest claims? A deeper problem with this proposal is that the Relevance View is supposed to be consistent with our living autonomous lives. It is supposed to apply to situations that are rare enough that we would not be constantly burdened by it and have to frequently appeal to a different domain of reasons to violate it.

Voorhoeve (2016) makes a more promising suggestion with regard to accommodating people’s preferences for risk-taking in healthcare settings. The public may be willing to contribute additional funding for healthcare that it would not otherwise provide had the healthcare sector distributed resources with strict adherence to the Relevance View. If that is so, it is incorrect to frame the resources that go into treating minor ailments as competing with the resources that are distributed in accordance with the Relevance View towards major health burdens. Should the public wish to make use of the efficiencies of a public health system to provide treatment of minor ailments, it may be prepared to put additional funds (in the form of paying more tax) towards the treatment of these minor ailments, rather than paying for these services out-of-pocket (Voorhoeve 2016, p.154).⁴¹ Voorhoeve’s suggestion provides a less arbitrary way to justify the resourcing of treatments for minor ailments because funding is essentially tied to different domains of health-related claims.

Despite these justifications for side-stepping the Relevance View in public healthcare settings, Voorhoeve (2018) cites the 1990 Oregon healthcare rationing exercise as an example where public sentiments seem to have actually aligned with the Relevance View:

Infamously, on grounds of cost-effectiveness, officials recommended that tooth capping should take precedence over treatment for terminal appendicitis (Ubel et al., 1996). This proposal generated public outrage, precisely as one would predict if people’s moral sentiments conformed to ARC [the Relevance View]. (Voorhoeve 2018, p.154)

Voorhoeve notes that in this case, the main beneficiaries (low-income individuals) were not the primary funders (state and federal government) and so the justification for the treatment of minor ailments due to funding by self-interested parties was not available. Therefore, this healthcare system is arbitrating between genuine competing claims to aid.

I conclude this section by restating two additional reasons that the Relevance View might not run against the practice of attending to minor ailments in the healthcare system. First, minor ailments

⁴⁰ See Herlitz (2018) for one such proposal that builds on the ideas of Chang (2013) who suggests that when ‘given’ reasons run out, societies can create their own reasons.

⁴¹ Taurek (1977) makes a related point based on collective ownership with regards to Anti-Aggregation: where the patients have some collective ownership of a resource, saving the many could be justified by agreeing in advance that the resource should be used to save as many people as possible.

evolve into more serious conditions if left untreated. Second, claims to the treatment of minor ailments may not compete with claims to the treatment of major conditions, if the former are mandated on different grounds to the latter.

5.2 Competing claims and risk

In the previous chapter, I argued that when all claimants have the same *ex ante* interests at a time, then this can be seen as within the domain of intrapersonal aggregation and so lies outside of the scope of the Relevance View. In social risk cases, we can choose for all as we would choose for them each individually (Frick 2015). Rather than this generating an inter-personal trade-off, we are in effect instead making a self-interested gamble on the possible futures of the same individuals (Thompson 2018). I also argued that when the *ex ante* interests of members of different groups conflict, we should aggregate on the basis of *ex post* claims (Walén forthcoming). This avoids awkward situations where *ex ante* reasoning seems to lead us to saving an identified life over any number of statistical lives (see §4.4.2). However, there are various practices where the *ex ante* interests of members of different groups conflict and aggregating on the basis of *ex post* claims would seem to lead to equally troubling results. I outline the problem in §5.2.1. In §5.2.2 I offer a solution.

5.2.1 A problem for competing claims that involve risk

Consider again Scanlon's adaptation of the *World Cup* case:

Suppose, for example, that we are deciding whether to build a new system of transmitting towers that will improve the quality of reception for many television viewers. It may be highly probable that in the course of this project a number of workers will suffer harms at least as great as Jones's. Yet we do not think that it is therefore wrong to go ahead. (Scanlon 1998, p.236)

In this adapted case, the *ex ante* interests of the television viewers conflict with the interests of the workers. Because it is not in the interest of any given worker to waive their *ex post* claim, and the viewers' claims are not relevant to these claims, it is plausible that we should not allow the transmitting towers to be built. However, as Scanlon suggests, so long as reasonable precautions are taken, the workers' informed consent seems decisive in this matter. By consenting to this project, each worker is essentially waiving their *ex post* claim.

However, situations with this structure come up in the real world where it is not reasonable to assume that we have the consent of those exposed to risk, nor would attaining their consent be viable. Consider the practice of convenience driving (Walén forthcoming); that is, driving that has no wider social importance other than it is convenient for those who engage in it. Individual convenience drivers, in one trip or even over the course of their lives, impose infinitesimal risks of very significant harms on people who choose not to engage in this practice; for ease let us call them pedestrians (It is important that convenience drivers and pedestrians are not the same people, otherwise this would simply be a case for intrapersonal aggregation.) We know pedestrians die every year, and at least some

of those pedestrian deaths are due to convenience driving. Therefore, pedestrians would have an interest in banning the practice of convenience driving. Because the *ex ante* interests of individual convenience drivers and individual pedestrians conflict, we should aggregate on the basis of their *ex post* claims (Walen forthcoming). Plausibly, the claims of convenience drivers are much weaker than the *ex post* claims against severe driving-related injury or death, and the Relevance View would rule in favour of pedestrians.⁴²

However, banning convenience driving for the sake of a few pedestrians seems counter-intuitive. One reason this is counter-intuitive is because if there were only one pedestrian it is not reasonable to assume that they have a full *ex post* claim against death. If there existed only a single pedestrian in the population, their *ex post* claim should surely reflect the fact that they are very unlikely to actually die. It still would not be in this person's *ex ante* interest to waive their *ex post* claim but this discounted *ex post* claim may be relevant to the claims of individual convenience drivers in which case it can be outweighed. It is the mere fact that there are enough pedestrians to generate one undiscounted *ex post* claim against death that means the Relevance View seems to ban this practice. Even if you do not find this result troubling, consider another case such as eating peanuts.⁴³ It would be strange if there should be a general prohibition on the consumption of peanuts just because it is more than likely that someone with a severe peanut allergy will die from accidental exposure. Again, this is not to say that reasonable precautions regarding their consumption should not be taken. It is plausible that restaurants should still make an effort to give people with peanut allergies reliable information and handle their allergy with care. But banning the consumption of peanuts just because across an entire population we can expect one peanut-related death is deeply counter-intuitive. I will now show why the Relevance View may not necessarily condemn such practices as convenience driving and the consumption of peanuts.

5.2.2 Differentiating between acts and conjunctions of acts

Tenenbaum discusses a similar set of problems currently facing the Relevance View in the criminal justice domain:

A system of criminal punishment will over time almost certainly condemn some innocent people. Still, it seems plausible to say that despite the serious harm the justice system will

⁴² The *ex post* claims of pedestrians may be even stronger along another dimension: by allowing convenience driving, the transport agency is imposing risk on pedestrians in order to benefit convenience drivers (Walen forthcoming). They are harming in order to benefit rather than the harm coming as a side-effect of benefiting convenience drivers. It is a common deontological conviction that this is harder to justify and one reason this may be the case is that we have stronger claims against being harmed as a means rather than as a side-effect.

⁴³ Ashford (2003) uses Scanlon's example of air travel: there is a small chance that an airplane will fall out of the sky and kill someone on a remote island who does not benefit from air travel. Ashford expresses this particular problem with risk as a *demandingness* problem: giving up air travel would be too demanding on us. Ironically, there is a real-world problem frighteningly close to Scanlon's example which does seem objectionable: the disproportionate effects that airplane emissions have on people being displaced due to sea-level rise.

cause to certain people, there is no general deontological prohibition against a system of criminal punishment. However, things would be different if the risk were concentrated in one person and in one action. We would not replace our system with one such that we would punish Jane who is (almost certainly) innocent but that would somehow allow us to reduce the number of people who are wrongly condemned overall. (Tenenbaum 2017, p.700)

We can accept a general deontological prohibition against imprisoning innocent people. Still, we would not take this to prohibit locking anyone away when we are not absolutely certain of their committing the crime. Of course, an agency may legitimately take it upon itself to minimise the number of innocent convictions. However, what seems to matter is that each sentence be justifiable to each convicted person. While it makes sense to look at the permissibility of the set of these acts as a whole, it is also meaningful to view them as individual acts. It is easy to view this argument in terms of mathematical probabilities. If a reasonably reliable judge passes a guilty verdict with 99% confidence, after 100 such verdicts, we can reasonably expect her to have judged at least one case wrong. But Tenenbaum (2017) wants to say that the judge may be fully justified in every verdict she passes, and we can still say that the criminal justice system may have certain unjustified features. As Tenenbaum puts it, “the deontic status of a pattern of activity might not supervene on its constituent acts” (p.706).

We can make a similar case for convenience driving (and peanut consumption). Once convenience driving is understood as a series of individual acts rather than a single complex act, pedestrians’ *ex post* claims no longer combine into a full *ex post* claim against being killed. This means that the *ex post* claims of pedestrians may now be relevant to the claims of convenience drivers because each individual’s *ex post* claim should reflect the low probability of severe injury resulting. It is plausible that these claims are now relevant to each other.

This discussion makes sense of a much simpler idea in deontology. In the context of imposing risk of death, Tenenbaum (2017, p.684) writes that “[i]t seems perfectly coherent to say that it is disrespectful to seriously endanger someone’s life, but that we do not owe it to him or her to guarantee his or her safety”. It is sometimes inappropriate to view the conjunction of convenience drivers as a single agent, and to view this single agent as imposing a risk on pedestrians that is equivalent to seriously endangering them (given that some of them are likely to die). This means that while we can still set policy goals to reduce the harm that convenience drivers impose on pedestrians, the Relevance View does not rule out this activity on the grounds that a full *ex post* claim can be conjured together.⁴⁴ The correct scope of competing claims is between individual actors rather than between groups.

⁴⁴ There may also be a place for compensation in this debate. Firstly, claims may give rise to secondary duties of compensation. This applies generally for satisfying competing claims. If Doctor Debbie provides Barb and Claire with the medicine rather than Anne, for example, perhaps Debbie or Barb and Claire owe it to Anne that she be benefited in another way. In the context of risk, people may be willing to waive their *ex post* claims if a reasonable compensation policy is implemented. For example, pedestrians may be willing to waive their *ex post* claims if they are generously compensated

5.3 Conclusion

In this chapter, I have aimed to highlight some ways that the Relevance View might be less revisionary to our current practices and policies than it initially seems. In healthcare settings, as well as highlighting some of the resources I had mentioned in previous chapter I brought further resources to bear on this problem. I argued that properly acknowledging the impact of harms on one's life as whole may justify the treatment of minor ailments. I also argued that the healthcare system may be accommodating people's preferences for risk in numerous ways.

In the second half of this chapter, I offered a rationale for defending certain practices where one group impose small chances of significant burdens on another group. Although we should use the *ex post* model when different groups' interests compete, we should not arbitrarily group people's *ex post* claims together. Sometimes certain policies are better thought of as arbitrating between individual actors, rather than between big groups.

The discussion in the next chapter adds an additional layer of complication to prioritisation in healthcare settings. Agencies making rationing decisions are often faced with a spectrum of claims of different strength. This raises the question of how the Relevance View can accommodate three-way choices and choices involving competing non-homogenous sets of claims.

should they be severely injured. We could actually view such a compensation policy as reducing pedestrians' *ex post* claims. One would have less of a complaint if one suffered serious bodily injury if one were given generous monetary or healthcare compensation. Of course, there are limits to this reasoning. We may not be able, even partially, to compensate for the loss of certain goods. Along these lines, for example, it is not clear how one could compensate for someone's death.

Chapter 6

Developing the Relevance View

In the real world, a single policy decision will affect different groups of people in different ways. In this chapter, I move beyond cases that resemble the *World Cup* and *Life versus Paraplegia* cases. We can call both of these one-many cases. In the *World Cup* case, one strong claim to aid competes with many very weak claims. In the *Life versus Paraplegia* case, one strong claim competes with many moderate claims. What was said in the one-many cases would also hold for many-many cases. If Jones's colleague too had an arm trapped in the electrical cabling, then because one strong claim is enough to outweigh the claims of the fans, two strong claims will also be enough. On the other hand, if a larger number of people's lives are threatened, then because claims against paraplegia are relevant to claims against death, we might expect this larger number of claims against death to outweigh the 1,000 claims against paraplegia. What remains distinctive about these cases is that (1) the competing sets of claims are homogenous, and (2) there are just two competing sets of claims.

In many settings involving competing claims, (1) and (2) are unlikely to obtain. For example, in the provision of publicly funded healthcare resources, many different types of claims compete for the same limited funding. I begin this chapter by introducing three-way choices, where it is not clear how the Relevance View guides us to choose. I show that the problem of such three-way choices has other analogies in ethics. Nonetheless, it is not clear how these problems should be resolved. One way to generate potential resolutions to these problems is to find general principles that are intuitively plausible in other situations and have promise here too.

The second part of this chapter involves looking at two-way choices that involve heterogeneous claims. There has been some progress towards meeting challenges of this sort which I will outline in §6.2. However, rather than try to develop these solutions further, I suggest an alternative solution that involves limiting the ambitions of the Relevance View.

6.1 Three-way choices and transitivity

In Chapter 2, I introduced two kinds of defeating relations. The first is the outweighing relation. For example, the 1,000 claims against paraplegia outweigh the single claim against death. The second is the disabling relation. This obtains when a sufficiently strong claim prevents a set of weaker claims from grounding a duty. For example, Jones's claim disables the claims of the football fans from outweighing it.

To introduce the next set of problems, I will use some shorthand. I will be comparing weak, moderate and strong claims. The kind of claim will be shown in parentheses, (). The number of people with this

claim will be shown using the number in front of the parentheses. A bar “|” will be used to signal that the respective claims are in competition (that is, they cannot be jointly served). I will “**bold**” the set of claims that the development of the Relevance View that I am entertaining tells us to satisfy. The cases I introduce might be thought of as analogous to situations where a healthcare agency must decide which kinds of harms to invest limited funds in preventing. In the following cases I just stipulate how many claims are sufficient to outweigh another claim and which claims are relevant to each other. This is simply to illustrate certain problems.

First, consider a case involving an outweighing relation:

1(strong) | **20(moderate)**

In this case, the agent can either serve the single strong claim or 20 moderate claims. Suppose that moderate claims are relevant to strong claims, and that 20 moderate claims is sufficient to outweigh the single strong claim. Therefore, according to the Relevance View, we should satisfy the 20 moderate claims. Consider another case involving an outweighing relation:

20(moderate) | **1,000,000(weak)**

In this case, the weak claims are relevant to the moderate claims. The 1,000,000 weak claims are also sufficient to outweigh the 20 moderate claims. Therefore, we should satisfy the group of weak claims. Finally, consider the following case involving the disabling relation:

1,000,000(weak) | **1(strong)**

In this case, the agent can either satisfy the single strong claim or the 1,000,000 weak claims. By stipulation, weak claims are not relevant to strong claims, so we are required to satisfy the strong claim.

The Relevance View raises some interesting questions about transitivity. Consider the following choice where we can only save one of the three groups:

1(strong) | 1,000(moderate) | 1,000,000(weak)

Here is a situation where it looks like the moderate claims defeat the strong claim, the weak claims defeat the moderate claims, and the strong claim defeats the weak claims. On the face of it, the relations among these competing claims seem to violate transitivity because it implies that what is true of the relation between A and B, and B and C, is not true of the relation between A and C. If this did violate transitivity one worry would be that it looks like whichever option we choose, we might have acted wrongly (Parfit 2011, p.203). For example, if we satisfy the strongest claim, we have acted wrongly because the moderate claims outweigh this claim. However, the relations between these items do not necessarily violate transitivity once we probe the Relevance View some more.

6.1.1 Temkin on transitivity

Temkin (2012, pp.195-7) discusses the transitivity of the “permissibility” relation explored by Frances Kamm.⁴⁵ It may be permissible for me to go to a movie rather than incur a significant risk to save a stranger’s life. It may be permissible for me to save a stranger’s life rather than keep a promise I had made earlier. And yet, it would be impermissible for me to go to a movie rather than keep a promise. The supposed intransitivity disappears when all three options are available because going to the movie is no longer an option. The permission to go to the movie only holds in the case where I have not made the promise. One similarity that we might observe is that we should only satisfy the weak claims when they do not compete with a strong claim. Because weak claims do compete with strong claims in the three-way choice above, these weak claims should be eliminated from the competition. However, opting for this response would be too fast, as there are other ways of specifying the Relevance View.

Temkin (2012) discusses an example concerning affirmative action policies in job hiring settings, which has striking similarities to the above three-way competing claims set-up. Suppose that an employer ought to offer the job to the candidate with the highest credentials. However, there is an exception to this rule. Owing to the historical relationship between African Americans and White Americans, when an African American’s credentials are close enough to those of a White American, the employer ought to give the job to the African American.⁴⁶ Now suppose that three candidates apply for a job and can be ranked in terms of their credentials in the following order: White American, Mexican American, African American. Given the choice between the White American and the Mexican American, the employer ought to offer the job to the White American. Given the choice between the Mexican American and the African American, the employer ought to offer the job to the Mexican American. Given the choice between the White American and the African American, the employer ought to offer the job to the African American. However, when considering the three options together, it is not immediately clear what the employer should do:

[I]f we decide that the correct policy of affirmative action “merely” generates the conclusion that it would always be wrong to have a white occupy a position for which a sufficiently qualified African American applied, then [...] we ought to hire Mexican American. However, we may decide that the correct policy of affirmative action “merely” tells us to rank all candidates in terms of their job-related characteristics and then choose the top one, unless the top one is White, and the second one is a sufficiently qualified African American, in which case choose the second one. This view would support [hiring the] White. Finally, we might decide that the correct policy of affirmative action tells us

⁴⁵ When Kamm (2007, p.298) talks about the Relevance View she does so in the language of respect and therefore as an implication of the Relevance View concludes that “disrespectful to consider” is not transitive. However, Temkin (2012, p.196) notes that a set of relations does not really violate transitivity when the features that lead us to choose A over B and B over C are not the same features that lead us to choose A over C. This is the case in our example, where two kinds of defeating relations are at work: outweighing and disabling.

⁴⁶ Of course, including only these two factors are intentionally very narrow. Temkin does not take a position on affirmative action policy itself, he just points out that this example has a relevantly similar structure to the Relevance View. For example, diversity in the workplace surely also counts towards why we might choose a candidate of one social group over another.

to rank all candidates in terms of their job-related characteristics and then choose the top one, unless the top one is white, in which case we should hire the top-ranked African American candidate, as long as he is sufficiently qualified relative to the top-ranked white. This would support [hiring the] African American. (Temkin 2012, p.200)

Notice that the problem only occurs because the affirmative action policy is crucially underspecified. I will not try to specify the correct affirmative action policy, but I will lay out the respective options for the Relevance View as Temkin does in the job hiring case.

We might decide that whenever there is a set of claims that is not relevant to another set of claims in a competition, we should eliminate this first set of claims, and settle the remaining competition by seeing which set of claims is stronger.⁴⁷ In our example, because the strong claim disables the weak claims, we can immediately eliminate the weak claims from the choice set. By stipulation, in the two-way choice between satisfying a strong claim or 20 moderate claims, we should satisfy the moderate claims. Therefore, according to this justification, we should satisfy the moderate claims.

Alternatively, we might decide that when one set of claims is relevant to another set of claims in the choice set, the disabling defeater is neutralised, and we can decide on the basis of the weighing relation.⁴⁸ In our example, because the moderate claims outweigh the strong claim, and because the weak claims outweigh the moderate claims, we should satisfy the weak claims.

The final alternative might simply have us defeat in favour of the strongest claim.⁴⁹ In our example, we can defeat in favour of the strong claim using the following procedure: the weak claims outweigh the moderate claims, and the strong claim disables the weak claims, so we have a justification for satisfying the single strong claim.

Rather than privileging one of these options, one could settle on the position that these options are equally good, or on par, and the only rule governing our choice is that we save one group (Temkin 2012, p.201). However, another way to settle which option is most promising is by looking at complex

⁴⁷ Voorhoeve (2014) argues for this solution.

⁴⁸ Tadros (2019) supports this solution but suggests that other resources in deontology complicate this picture. For any option to win the competition, it must rely on a second option to defeat the third option. For example, if we have a duty to satisfy the weak claims, then this is because the moderate claims outweigh the strongest claim. But Tadros suggests that this introduces a morally important consideration: the people with moderate claim may need to be strengthened to reflect the fact that they are being used to benefit the weaker claimants.

⁴⁹ This option might be supported by the suggestion due to Van Gils and Tomlin (forthcoming) that one could treat the defeating process as always defeating in favour of the strongest claim. Van Gils and Tomlin suggest a procedure which tells us that when there are many ways one could match claims overall, we should match in favour of the overall strongest claim (pp.31-3). They call this procedure “strongest decides”. However, they reject this procedure because it could lead to results whereby the addition of a claim to the competition could actually prevent the group with which this additional claim is aligned with from having its claims satisfied. (Illustrating this here would take up too much space.)

two-way choices. We can eliminate certain options if they generate counter-intuitive results in certain set-ups. I will explore this strategy now.

6.2 Complex two-way choice cases

The previous section puts us in a position to begin investigating how the Relevance View can guide us in complex two-way choices, using cases which involve potential outweighing and disabling relations. Patrick Tomlin (2017) identified a strange implication of the Relevance View when it comes to certain types of these cases: the group of claims that one ought to serve switches in counter-intuitive ways when other claims are added to the competition. Using my shorthand method again, I will “underline” claims below to show that they have been added to the competition. Take first the view called “Anchor by Competition”, which tells us to eliminate claims that are not relevant to at least one of the claims with which they compete (Horton 2018). While this might initially seem sensible, it leads to strange results. Suppose that in the following case we are required to serve the claims on the right-hand-side (RHS):

(A) 1(strong) | 20(moderate)

Observe what would happen if weak claims are added to the competition in the following way:

(B) 1(strong) + 1(weak) | 20(moderate) + 1 billion(weak)

Even though many more weak claims were added to the RHS than to the left-hand-side (LHS), the single weak claim added to the LHS was enough to tip the balance in favour of the LHS. This is because, whereas the weak claim added to the LHS is relevant to the moderate claims with which it competes, the weak claims added to the RHS are not relevant to the strong claim with which they compete. This means that the billion weak claims are disabled and the single weak claim is not. This counter-intuitive switch is enough to reject Anchor by Competition.⁵⁰

A different view is “Anchor by Strength”, which tells us to eliminate claims that are not relevant to at least one claim in the competition. In case (B) above, the weak claims added to both sides would be eliminated because neither of them is relevant to the strong claim in the competition. Therefore, we should choose in (B) as we would choose in (A). However, Anchor by Strength leads to another kind of counter-intuitive switch:

(C) 4000(weak) | 20(moderate)

(D) 4000(weak) + 1(strong) | 20(moderate)

⁵⁰ In their papers, Tomlin and Horton also develop plausible principles which these different developments of the Relevance View violate, to show where these developments go wrong. Rather than restating these principles here, I take the counter-intuitive results enough to warrant rejection.

In (C), the weak claims are enough to outweigh the moderate claims, so we are required to serve the LHS. However, the addition of a strong claim in (D) now disables the weak claims with which it is aligned. Because the moderate claims are enough to outweigh the strong claim alone, we are now required to serve the RHS.

Both Anchor by Competition and Anchor by Strength are called Global Relevance views because claims are either relevant to the competition or they are not relevant to the competition. Local Relevance was proposed in response to these problems by Victor Tadros (2019). Local Relevance tells us that all claims have the potential to outweigh competing claims to which they are relevant. Local Relevance has problems of its own, but these problems are at least more palatable.

(E) 1(strong) | 4000(weak)

(F) 1(strong) + 20(moderate) | 4000(weak) + 20(moderate)

In (E), the strong claim disables the weak claims, and so we have a duty to satisfy the strong claim. In (F) the moderate claims added to the RHS are enough to outweigh either the strong claim or cancel the moderate claim added to the LHS. Allowing the moderate claims to cancel each other would leave us in the same situation as (E). However, we can avoid invoking the disabling relation by having the moderate claims on the RHS outweigh the strong claim and having the weak claims outweigh the remaining moderate claims. It might seem odd that the verdict of (E) differs from that of (F), even though the claims added to either side in (F) are the same. However, Tadros (2019, p.181) argues that this is not so implausible. Adding the moderate claims to the first case gives “defeating power” back to, or “activates”, the weak claims in virtue of being relevant to moderate claims. Because the weak claims are now relevant, they can do the outweighing work that they could have done in (E) had they not been disabled.⁵¹

6.2.1 Kamm on irrelevant utilities

Frances Kamm (1998, p.101) has advanced another kind of example that has been used to support the Relevance View. This example poses another challenge to Local Relevance. Kamm suggests that when an agent can only save one of two lives, the fact that saving one of those lives will also allow the agent to prevent the headache of another person should not be able to tip the balance in favour of saving that life. Kamm thinks that determining which life should be saved on the basis of a single headache would be *unfair* to the person whose claim is thwarted. These cases, Kamm suggests,

⁵¹ There is another question about whether irrelevant claims can break ‘almost-ties’. Brown (2019, p.16) puts it well when he suggests that “[i]f one allows irrelevant harms to break ties, then there seems little justification for not allowing them also to break almost-ties, especially given that an almost-tie can be arbitrarily close to a tie.” Tadros (2019, p.178) is also favourable to this option. At first this move seems *ad hoc*. However, the contribution of minor claims could be capped. Weak claims cannot defeat strong claims, but weak claims and moderate claims can defeat strong claims.

demand a coin toss. To use the language introduced earlier, claims against headaches are just not relevant to claims against death.

Kamm's example is inconsistent with Local Relevance because once we are at a tie with respect to the two death, the single headache because locally relevant and so is able to break the tie. However, one reason the Kamm's example seems compelling is that a single headache is within the margin of error for what is at stake between the two individuals that face the prospect of death. It is incredibly unlikely that the only difference between what is at stake for the two death-threatened individuals is equivalent to a mere headache. For example, one individual may have completed a major life project while the other is yet to achieve one.

However, one problem with Kamm's example is that even if a single headache could not tip the balance, it is plausible that a large number of headaches could. Suppose, for example, that we could fund a drug that would save the life of A or we could fund a cheaper drug that could save the life of B and because it was cheaper, we could also fund a year of paracetamol to a million people. Surely, it would not be unfair to A to choose the second option.

I think the correct way to accommodate the intuition from Kamm's example with Local Relevance, is to concede that fairness has some weight in these cases without giving up the idea that headaches should count for something (Hirose 2015; Tadros 2019). All we need to do is be open to pluralism about what is important when claims compete. Satisfying the strongest set of claims is one consideration, but fairness is another consideration. In Kamm's example, perhaps fairness is the dominating consideration because the difference between what is stake between the two individuals is in the same ball park. However, when we consider one million people's claims to have their headache ameliorated becoming aligned with one of the death-threatened individuals, for example, it becomes less plausible that fairness is the decisive consideration.

I conclude this section by reiterating that Local Relevance is a promising development of the Relevance View.⁵² Moreover, Kamm's tie-breaking argument should not be used to motivate the Relevance View as it appeals to an external consideration.

6.3 Conclusion

In this chapter I have explored how the Relevance View can be extended beyond homogenous one-many cases and the various challenges that this presents. In §6.1 I considered three-way choices that seem to give rise to intransitive relations. I argued that this intransitivity is not obviously troubling. First, these kinds of situations come up elsewhere in ethics. Second, resolving the tension is a matter of specifying the view, not a problem with the view itself. In §6.2 I considered complex two-way

⁵² I note here that Aart Van Gils and Patrick Tomlin have a paper coming out where they identify several ambiguities with Local Relevance. These problems arise out of trying to specify Local Relevance further to handle larger groups.

choices. Local Relevance looks like the best refinement of the Relevance View. Moreover, Kamm's tie-breaking argument should not be used to motivate the Relevance View as it appeals to the external consideration of fairness.

Local Relevance will have some bearing on worries about prioritisation presented in the last chapter. Perhaps we need not appeal to an independent moral mandate to accommodate the treatment of minor ailments. Instead the fact that competing claims are come in a spectrum in public healthcare settings means that minor claims will often be justified because they are 'locally' relevant to more significant claims.⁵³

This chapter concludes the five substantive chapters of this thesis. I have aimed to limit the scope of the Relevance View in numerous ways for both theoretical and practical reasons. All of these restrictions make the Relevance View less revisionary. But does this mean that the Relevance View has no real-world import? This is one of the questions I will explore in the conclusion.

⁵³ There remain further complications with this proposal which I outline in f.n.49.

Chapter 7

Conclusion

The aim of this thesis is to offer a general defence of the Relevance View. The Relevance View tells us when claims of one strength can outweigh claims of another strength. When claims compete, the fact that one kind of claim is sufficiently weaker than a competing kind of claim means that the weaker claims, regardless of their number, cannot outweigh the stronger claim. The thought is that these weaker claims are not relevant considerations when they compete with such strong claims. In this, the final chapter, I will first offer a summary of the arguments presented and conclusions reached in Chapters 2–6. I will then draw some general conclusions from my discussion and point out some areas for future research.

7.1 Chapter summaries

Chapters 2 and 3 of this thesis were centrally concerned the theoretical basis of the Relevance View. The goal of Chapter 2 was to establish the best rationale for the Relevance View. While the Relevance View can accommodate two important intuitions, without something further to be said in its favour it would appear *ad hoc*. I followed Nagel in arguing that Pure Aggregation and Anti-Aggregation are appealing in the divergent ways they understand equal moral concern. Pure Aggregation understands equal moral concern in terms of claims adding and subtracting in their contributions to the moral quality of an act. Anti-Aggregation treats equal moral concern as making pairwise comparisons, staying close to each person's own viewpoint, and satisfying the strongest individual claim.

I next differentiated the Relevance View from other views such as the Threshold View, Maximin, and Scanlon's tie-breaking argument. This opened the door for me to defend the Relevance View within a deontological framework. The last task of Chapter 2 was to find an independent rationale for the Relevance View. I rejected Voorhoeve's test for two reasons: one-one duties do not obviously translate into one-many duties, and the idea of preferring that another person's claim be served suffered from a tracking problem. Finally, I argued that the Relevance View was best understood as a property of Weak Aggregation (Walen forthcoming). As a burden is distributed, the claims that it grounds aggregate less and less strongly. I also explored a value that might underlie Weak Aggregation. The value of solidarity seemed important. When a burden is distributed, the many are "sharing the weight" and thereby "taking the load off another's shoulders."

In Chapter 3, I responded to sceptical arguments against the intuitions that motivate the Relevance View. These arguments aimed to show that we are imagining the *World Cup* and *Life versus Paraplegia* cases incorrectly and they are thereby eliciting the wrong intuitions. I argued contrary to these arguments that we have good reason to trust our intuitions in these cases. Of course, our intuitions could be wrong, but a more convincing full-blooded comparative endeavour would be

needed to show why this is the case. One alternative to accepting the truth of the Relevance View was that the Relevance View was justified on grounds on moral uncertainty: we should use Pure Aggregation and Anti-Aggregation in different cases as a way of minimising how wrongly we act if it turns out that one of these rationales is false. However, the Relevance View has parsimony in its favour. Neither does it merely gain credibility by sharing characteristics with these two views. As I argued, the Relevance View does have its own unique rationale.

The remaining substantive chapters of this thesis were united in that, taken together, they were a general attempt to limit the counter-intuitive or revisionary implications of the Relevance View. In Chapter 4, I outlined two general reasons to limit the scope of the Relevance View by appealing to intrapersonal considerations. It is plausible that we should not be guided by the Relevance View when doing so would make everyone worse off. There is another plausible intrapersonal justification for taking certain kinds of social risks. I argued that when it is in everyone's *ex ante* interest (at each moment in time) then we should take such risks. Indeed, any theory of beneficence must tolerate some risk-taking on the part of others in order to preserve the congruence with an individual's pursuit of a good life (one where they are not impoverished by excessive risk-aversion). We should nonetheless be cautious when using *ex ante* reasoning. We cannot just adopt any policy that will be in everyone's *ex ante* interests. I argued that we should adopt Frick's decomposition test for assessing whether a risk is permissible.

In Chapter 4, I also considered two objections to my restricted use of *ex ante* reasoning. First was the objection from dependent risk. I concluded that while dependent risk looked more problematic than independent risk, it is not problematic enough to threaten the permissibility of taking social risks. The second objection was that *ex ante* reasoning gives unconstrained priority to the identified life in certain situations. I argued that the *ex ante* reasoning only applies when individuals' interests are aligned. When they instead compete, as they do in the identified versus statistical life cases, we should revert back to the *ex post* reasoning (Walen forthcoming).

In Chapter 5, I aimed to show that the Relevance View would not be implausibly revisionary in certain settings. First, the consideration of self-interest may generate independent reasons to treat minor ailments over more severe diseases. Following Voorhoeve (2017), we could see the provision of services for minor ailments as coming from additional funding that would not otherwise be available. The second area where the Relevance View looked implausibly revisionary is where one group imposes a small risk of a significant harm on another group for the sake of a relatively trivial benefit. I used the examples of convenience driving and the consumption of peanuts. I argued that we cannot always just conjure together a situation that looks like a violation of the Relevance View. We must have good reason to think that the acts in question agglomerate meaningfully.

In the final substantive chapter of this thesis, I looked at different proposals to extend the Relevance View beyond conflicts between two homogenous groups. I argued that non-consequentialism is well

placed to handle the supposed intransitivity that arises from the view in such settings. The problem of resolving intransitive sets of competing claims was one of disambiguating the Relevance View and not that the view is self-defeating. I also agreed with the current state of the literature that Local Relevance (Tadros 2019) is the best development of the Relevance View.

7.2 General conclusions

One of my overall strategies in this thesis was to bring numerous resources from deontology to bear on the Relevance View. In Chapter 2 I noted that there are at least four reasons to favour deontology that should now be even more apparent at this end of the thesis. The first was that violating the Relevance View displays a kind of disrespect towards the person with the strongest claim, and this places a constraint on an agent that is independent of whether or not acting in this way brings about the best state of affairs. A second reason was that the transitivity issues outlined in Chapter 6 are not as problematic for the deontologist as they are for the consequentialist. A third reason to favour deontology is that it makes sense of the notion of waiving one's claim that I made use of in Chapters 4 and 5 to explain why agents may be able to act in people's *ex ante* interests even though this would give rise to a pattern of *ex post* claims that seem to violate the Relevance View. A fourth reason to favour deontology concerned the structure of two kinds of defeating relation that show up elsewhere in deontology: disabling and outweighing. Sometimes different considerations seem to pull against each other and the relative weight of each consideration is decisive. Other times different considerations seem to pull in opposite directions but one consideration strips another consideration of its normative force (e.g., a doctor operating on a patient without their consent). This follows the structure of the Relevance View.

Whether the Relevance View is theoretically superior to other competing views or whether it should form the basis of real-world decision making, is a topic for another day. I accept that for practical reasons mainstream theories of distributive justice—such as utilitarianism, prioritarianism, and egalitarianism—may be indispensable. Certain problems may be too complex for us to follow the Relevance View. As was made clear in Chapter 6, agencies would not only have to give every health-burden a weighting but also determine which other burdens it is relevant to. Future research should weigh these considerations—practical utility and theoretical truth—against one another.

In light of the Relevance View's potentially practical failings, I want to acknowledge that this there is new and challenging work coming out all the time that bears on the central themes of this thesis. One paper worth mentioning is by Van Gils and Tomlin (forthcoming). Van Gils and Tomlin (forthcoming) argued that Tadros's Local Aggregation is itself underspecified. In their paper they point out that some of the most promising specifications lead to further counter-intuitive results or yield their own ambiguities. Merely identifying ambiguities is not knock down argument for the Relevance View, it merely shows that it needs further development in certain complex settings. Nevertheless, engaging

with these ambiguities is definitely a task for defenders of the Relevance View to consider in the near future.⁵⁴

During the course of this research project, I discovered that many areas of public policy make assumptions about aggregation. In future research it would be interesting to see if justifications along the lines of the Relevance View are used or could valuably be used in these contexts. In the more dynamic domain of public debate, I think that a wider awareness of the Relevance View may help us overcome certain impasses. The Relevance View identifies a particular kind of complaint that needs to be taken seriously: “their claims are not relevant to mine!” Before we even begin thinking about aggregation, we need to ask whether the respective claims are relevant to one another. If one set of claims is not relevant, then there is no point in debating which factors should be included and how much weighting we need to give when determining which claims to serve. To use a controversial but illustrative example, we may not need to settle whether society’s combined pleasure from eating meat outweighs the suffering of certain animals if everyone gets on board with the idea that our pleasure may not be relevant to their suffering. Some parties to a moral decision may agree that some minor claims appear trivial in light of more pressing competing claims but because of a theoretical commitment to Pure Aggregation, they may feel uncomfortable taking a position on whether these trivial claims could outweigh the more pressing claim(s). The Relevance View undercuts this worry by valuing the initial intuition that some considerations should not be counted.

With this all said, I do not believe that something as abstract as the Relevance View is in the game of reforming our practices. However, to quote Hirose (2015): “there are various things we should expect out of ethics, and that clarifying the structure of ethical thinking is one of them” (p.222). I believe the Relevance View does just that. The Relevance View amounts to a theoretically plausible middle way between Pure Aggregation and Anti-Aggregation. It explains how we can hold onto two intuitions that initially seem incompatible. Moreover, if Weak Aggregation is the correct rationale for the Relevance View, then the Relevance View also brings our attention to how solidarity enters into our distributive justice judgements. This value terminates in the position where we collectively refuse to let burdens fall on some where smaller burdens can be borne by others.

⁵⁴ Another area of particularly relevant contemporary work is to do with the notion of decomposition in deontology. I relied on Frick’s decomposition test to constrain *ex ante* rules in Chapter 4. Gordon-Solmon (2019) has recently criticised Frick’s proposal on the grounds that it cannot make the morally important distinction between acts that involve two decisions and acts that make future decisions impossible. Both Frick and Gordon-Solmon are writing in the context of contractualism, so it is not clear exactly how the worries of Gordon-Solomon would translate to the Relevance View. Walen’s (forthcoming) wonderful paper that I have engaged with quite a bit throughout this thesis also discusses issues to do with decomposition. In that paper he responds to several arguments from Horton (2017) and Hare (2016) that find problems with using the *ex post* model for risk. Horton and Hare variously show that where acting one way in certain cases seems to violate a deontological requirement, these wrongful acts can decompose into smaller acts that are each morally required.

Appendix A

An Alternative Response to the Argument from Risk

In Chapter 4, I argued that the Relevance View is consistent with taking productive social risks. Respecting the Unity of Individuals requires that agents tolerate taking some risk on the part of others. Because social risks involve gambling on the possible futures of the same individuals, taking social risks are not bona fide interpersonal trade-offs and so lie outside the scope of the Relevance View.

However there is an alternative strategy that preserves the permissibility of taking risks on behalf of individuals that may be consistent with the Relevance View. Consider again the following stylised public health emergency from §4.3.1:

One million children have contracted a virus. If the public health agency does nothing, all one million children will be ill for one week and then die. The public health agency can produce one of two vaccines that must be taken by all one million children to be effective.

Vaccine 1 is certain to save the life of each child but will do nothing for their being ill for one week.

Vaccine 2 is certain to prevent each child from being ill for one week and save their life, apart from a 1 in 100,000 chance that it does nothing.⁵⁵

The Argument from Risk challenges the Relevance View by arguing that if, according to the Relevance View, selecting *Vaccine 2* in the multi-person case is impermissible, individual gambles are also impermissible. But individual gambles are permissible, so the Relevance View is false. Take the following case of an individual risk:

Joe has contracted a virus that will cause him to be ill for one week and then die; Joe's doctor has the choice of giving him just one of the two following vaccines:

Vaccine 1 is certain to save Joe's life but will do nothing for his being ill for a week.

Vaccine 2 is most likely to save Joe and relieve him of any illness, apart from a 1 in 100,000 chance that it does nothing.

The problem is that it seems perfectly permissible for Joe's doctor to choose *Vaccine 2* in this case. His chances of dying are so small that it is clearly in his personal interest to take this gamble rather than be ill for a week.⁵⁶ Any theory of beneficence should tolerate some level of risk taking on the part of others if it is not to be radically out of line with people's own reasonable risk preferences. However, there does not seem to be any important moral difference between the individual risk and the multi-person risk; the multi-person risk is just the individual risk taken by many children. So, we might

⁵⁵ Adapted from Dougherty (2013) and Frick (2015).

⁵⁶ Perhaps a certain amount of risk aversion is justified, but too much risk aversion is a sure way of ruining one's life. I like to imagine a hermit who will not leave the house and keeps the doors constantly locked.

conclude that selecting *Vaccine 2* in the multi-person case is also permissible, and if this involves an impermissible trade-off, then the Relevance View is false:

- (1) If the Relevance View is true, then selecting *Vaccine 2* in the multi-person case is not permissible.
- (2) If selecting *Vaccine 2* in the multi-person case is not permissible, then selecting *Vaccine 2* in the individual risk case is not permissible.
 - a. The multi-person case is just the conjunction of acts of the same sort as the individual risk case.
 - b. The permissibility of an act is the same as the permissibility of a conjunction of acts of the same sort.
- (3) Selecting *Vaccine 2* in the individual risk case is permissible.
- (4) Therefore, the Relevance View is false.

Premise (3) is hard to dispute. Not only is taking individual risks of this nature ubiquitous, but any plausible theory of self-interest would agree that *Vaccine 2* is better for Joe than *Vaccine 1* in the individual risk case. There would be an odd asymmetry if our theory of beneficence did not similarly tolerate taking some risk on the part of others. Therefore, the other two ways to resist the Argument from Risk are to reject premise (1) or to reject premise (2). To reject (1) is to claim that selecting *Vaccine 2* in the multi-person case does not violate the Relevance View because there is an important moral difference between choices made under uncertainty and choices made under certainty (such as the *World Cup* case). This is my preferred solution, for which I argue in Chapter 4. The other option is to reject (2). This has been a popular move in the literature. The easiest way to reject (2) is to show that the multi-person risk is not morally equivalent to the single-person risk. I will now show that various attempts to do this do not succeed. One general reason I think that this strategy does not succeed is that it fails to capture the distinctive inappropriateness of trading weak claims against much stronger claims.

A.1 Differentiating between social risk and individual risk

In the following sub-sections I will explore the second of the two ways to resist the Argument from Risk. This involves rejecting premise (2) which states that if selecting *Vaccine 2* in the multi-person case is not permissible, then selecting *Vaccine 2* in the individual risk case is not permissible.

There is a precedent in ethics (and particularly deontology) for differentiating between an act and a conjunction of acts of the same sort. What is distinctive about one person receiving *Vaccine 2* is that they themselves are not likely to die. In fact, any given person is expected to benefit greatly from receiving the vaccine. Everyone has the expected benefit of not being ill for a week minus the heavily discounted likelihood of death. However, after a certain number of people receive the vaccine, the chance that *someone* will die is exceptionally high.

Personal wrongness

Consider a case where I can use my boat to save some number of people at risk of drowning in the ocean.⁵⁷ For every additional person that I rescue, the chance that my boat will sink increases, and I know that if I rescue everyone my boat is almost guaranteed to sink. The single act of saving any one person may be morally required, but this does not mean that saving all of them is morally required. Indeed, there is some range of cases where I save incrementally more people, within which saving another person becomes impermissible.

Can an analogy be made with the vaccine case? Could we say that the nature of the conjunction of acts changes when (say) at least one death becomes sufficiently more likely? This move is unlikely to work. In the example of the rescue boat, the act of rescuing everybody is different from the act of saving just a few, because after some number of rescues every additional rescue threatens the lives of those already rescued. This is not what happens in the vaccine case. Everyone's risk is independent there; later acts do not threaten to make anyone vaccinated earlier worse off in any way.

For these reasons the direct appeal to personal wrongness will not work. Next I will consider impersonal wrongness.

Impersonal wrongness

Some egalitarians believe that inequality is impersonally bad; that is, at least part of what is wrong about bringing about unequal outcomes is independent of the concern for the worst off individuals. One might argue that this inequality is enough to show that the individual risk case is different from the social risk case. However, this response would miss the mark. We are interested in the distinct wrong that comes from violating the Relevance View's method of aggregating claims. Impersonal inequality may be a count against taking the social risk, but we are after a verdict that comes from the Relevance View itself. Nonetheless, we might be inspired by this appeal to impersonal wrongness and conclude that *Vaccine 2* is also impersonally wrong simply because of the pattern of harm that it brings about. However, this move will not work either. What is important about the Relevance View is the distinctive kind of disrespect we show to the person with the stronger claim when we trade irrelevant claims against that person's stronger claim. This distinctive kind of disrespect is not an impersonal value; instead, it is directed at the person with the stronger claim. This impersonally bad outcome could be objectionable; it could even make the conjunction of acts wrong all things considered, but this is not an implication of the Relevance View.

A role-based ethical idea

Deontology readily makes use of agent-relative duties. One simple example of this is special obligations that come out of close relationships. In the healthcare setting, doctors may have different

⁵⁷ This case is inspired by a case presented by Kirkpatrick (2018).

concerns than public health agencies. Whereas doctors may be centrally concerned with the health of their patients, public health agencies may be legitimately concerned with wider-population patterns. (That the duty of doctors is not directly derived from public health goals is evident, for example, in “conscientious objection”.)

Along these lines, we could say that the doctor *qua* doctor is justified in administering the *Vaccine 2*, but the state is not permitted to mandate the *Vaccine 2* programme because the scope of its considerations include the entire population. Perhaps the state’s power is overriding, and so the duty of Joe’s doctor is merely conditional—conditional on not being overridden by a higher authority. As I argued in Chapter 5, demarcating the spheres that the state may legitimately characterise as involving interpersonal trade-offs is important. But for the Argument from Risk currently under discussion, merely pointing to different authorities does not show why these authorities should view the situation differently.

More analogies with equality

If *Vaccine 2* is administered in the multi-person case, it is almost certain that some children will die and everyone else will escape incurring any burden. One way of justifying the *ex post* model would be to see if it makes sense of other relational values under risk, such as equality. Consider the case in which I give one of my children \$1. This seems like a nice thing to do. Now consider another case in which I give my other child \$100. This too seems like a nice thing to do. But it seems that I might do better not acting at all rather than giving one of my children \$1 and the other \$100 (Tadros 2019, p.198). Even though there is a sense that each of my children benefit from my conjunction of acts, there remains something objectionable about acting in this way.

The money case might seem objectionable because it may express favouritism to one of my children over the other. Consider the difference between giving the \$100 only to my blue-eyed child versus giving it to the child who chooses heads in a coin toss. The latter case looks less objectionable than the former because both my children have equal chances of receiving the \$100. However, what is distinctive about the case of social risk we are considering is that everyone has an equal chance of benefiting. Everyone has the same *ex ante* chances of benefiting from the vaccine. So, any complaint about favouritism is not available in the vaccine case.

Suppose then that I flip a coin to determine which of my children gets the \$100. Does my child who loses the coin toss (and receives the \$1) still have any grounds for complaint? There is a sense in which it still looks objectionable. Perhaps it is objectionable because it exposes their vulnerability to receiving benefits; receiving the \$1 instead of the \$100 is a significant loss, and it should not have been. Alternatively, determining someone’s fate by a coin toss fails to recognise the importance of benefits to them; it trivialises the importance of the benefit to them when in fact it has great importance to the individual. This indirectly disrespects the person. I could have given my children \$50.50 each, but I chose to gamble on their behalf. What this shows is that treating someone fairly is

not identical with giving them equal chances of benefit. It shows that treating someone fairly may be at least partially determined by the inequality of the outcome.

Can something analogous to the foregoing be said about the vaccine case? I think not. First, the odds look very different in the vaccine case. Ensuring that the outcome in the vaccine case is the same for everyone (i.e., choosing *Vaccine 1*) is not equivalent to everyone receiving \$50.50, but rather something closer to them each receiving \$5. The expected individual utility of *Vaccine 2* is 20 times better than the expected utility of *Vaccine 1*. Moreover, everyone's chances of receiving the equivalent of the \$100 are very good: 1 in 100,000. In other words, the cost of minimising *ex post* objections is huge.

What I take this brief discussion to have shown is that the procedure for selecting *Vaccine 2* is not obviously objectionable. By assumption, the decision shows no prejudice to certain groups. Moreover, the division of goods is not arbitrary: we are not trivialising certain interests because acting on the basis that some people will die comes at significant costs to everyone's *ex ante* interest.

I conclude that this general strategy of differentiating between the individual risk case and the multi-person risk case will not work. The better strategy is to accept that the multi-person risk case is permissible because it is in everyone's *ex ante* interest at the time the vaccine is administered.

Appendix B

Epistemic Considerations

In this appendix, I consider two epistemic assumptions that relate to the intrapersonal justifications I relied on in Chapters 4 and 5. First, I consider what epistemic considerations are important for an agent to act on the basis of *ex ante* justifications in social risk settings. Second, I consider an objection to *ex ante* reasoning which argues that such reasoning is epistemically irrational.

B.1 Knowledge and *ex ante* justifications

When each child has a 1 in 100,000 chance of dying by receiving *Vaccine 2*, this might be consistent with them having a 1 in 100,000 chance of having a particular gene that makes receiving *Vaccine 2* fatal. By assumption, the moral agent does not know which individuals have the fatal gene at the time of action, which is what makes *Vaccine 2* risky. But suppose that the agent could easily attain this knowledge; say, by administering a simple test. This generates two distinct questions:

- (1) Is the agent required to administer the test?
- (2) If the agent can administer the test, does this itself undermine the agent's justification for selecting *Vaccine 2*?

To answer these questions, it is helpful to consider the single-person case again. If any given child's doctor could run a simple test to determine whether or not he or she alone has the fatal genetic anomaly, this looks like something that would be in their interest. If any given child does not have the gene of interest, then there is no longer any risk for them in receiving *Vaccine 2*. If they do have the gene of interest, then we should give them *Vaccine 1* and they will be more than happy to spend a week being ill. Let us assume that there is a minor financial cost or inconvenience associated with administering the test to eliminate the risk, but this seems a worthwhile inconvenience to reduce one's risk of death from 1 in 100,000 to 0. From any given child's perspective, it is like looking twice before they cross the road; a minor cost for reducing their risk significantly. In the single-person case, the agent is required to administer the test. However, the mere fact that the agent can administer the test does not undermine their justification for selecting *Vaccine 2* because the patient may not have this gene.

If we rely on the permissibility of the individual case to justify the social risk case, then it seems we ought to administer the test before administering *Vaccine 2*. However, given that it is almost certain that some people will have the fatal gene in our population, we might conclude, without even running the test, that the mere fact that we can run the test means that we have reason not to administer

Vaccine 2. Because the fatal gene carriers are *knowable*, this gives the agent a duty to act against everyone's *ex ante* interest by administering *Vaccine 1* rather than *Vaccine 2*.⁵⁸

B.2 Epistemically irrational? A response to Fleurbaey and Voorhoeve

In their paper 'Decide as you would with full information!', Fleurbaey and Voorhoeve (2014) argue against using the *ex ante* model in multi-person risk cases. Their argument is most easily demonstrated using the dependent risk case from §4.4.1 where *Vaccine 1* that will leave every child ill is contrasted with *Vaccine 3* that is certain to be fatal for ten children but everyone else will avoid becoming ill. I have argued that if the agent knows the actual outcome of administering *Vaccine 3*, that is, knows who will in fact die, then administering that vaccine would be impermissible. The reason *Vaccine 3* is permissible is that no one knows which people will die.

The agent does know that administering *Vaccine 3* is impermissible in every possible world in which the agent knows the outcome (i.e., who will die). Fleurbaey and Voorhoeve (2014, p.120) argue that if an action would be impermissible in every possible world in which the agent knows the outcome, then it would be irrational to perform this action just because one is ignorant of which world will eventuate. However, contrary to Fleurbaey and Voorhoeve (2014), Frick (2014) argues that ignorance of the facts is morally important; it is the rationale behind *everyone* waiving their *ex post* complaint, and respecting this seems in many cases to be sufficient for defending the permissibility of the vaccine programme.

Frick (2014, pp.136-7) also points out a fallacy in the reasoning of Fleurbaey and Voorhoeve by making the following parody. If I knew which of two boats will win the boat race in advance, then I might decide not to watch the race. I know that in one world the red boat wins and in the other world the blue boat wins. But this knowledge of possible outcomes does not give me the same reason not to watch the boat race. My knowledge that *some* boat will win does not give me the same reasons as my knowing that a particular boat will win. Thus, agents are not compelled (as a universal principle of rationality) to act under ignorance the same way as they would act if they had full information.⁵⁹

⁵⁸ If there is not much cost associated with administering the test, then perhaps we should also administer the test but doing so is mostly redundant. Frick (2014) notes that the children would not want the agent to run the test because doing so would compromise their *ex ante* interest. I tend to think that consent is important in this case. If the children's caregivers will that *Vaccine 2* is administered, then perhaps we can do so. For a wider reaching discussion of how deontologists need to more honest about trading lives for money, see Fried (2018).

⁵⁹ Caspar Hare (2013, chapter 3) argues that there are some instances where it might be *irrational* to be guided by a fully informed agent. This argument relies on a controversial assumption about intransitive preferences, but it is worth mentioning. Suppose that I prefer A+ to A and B+ to B, but I have no preference of between A and B, A+ and B, and A and B+. Hare (2013) gives the example of having no preference when you can either save your Fabergé egg or your wedding album in a house fire. He notes that the fact that you know that there is a \$100 cash note next to the album in the right-hand room does not seem to make a difference to what you should prefer to save even though when given a choice between retrieving your wedding album or saving the wedding album and the \$100,

you prefer the latter. Now suppose that you keep the items in different rooms, but you do not know whether the album is in the left room or the right room. However, you do that the \$100 is in the right room. Hare suggests that even though your fully informed self who knows which item is in which room may have no preference for having one item retrieved over the other, you are not rationally compelled to act in accordance with your fully informed self and plausibly ought to enter the room on the right; where you have at least one reason to favour this room because it contains \$100.

References

- Ashford, E. (2003). 'The demandingness of Scanlon's contractualism', *Ethics*, 113(2), 273-302.
- Bramble, B. (2018). *The Passing of Temporal Well-Being*. Routledge.
- Cullity, G. (2018). *Concern, Respect, and Cooperation*. Oxford University Press.
- Dougherty, T. (2012). 'Aggregation, beneficence and chance', *J. Ethics & Soc. Phil.*, 7(2).
- Draper, K. (2002). 'The personal and impersonal dimensions of benevolence', *Noûs*, 36(2), 201-227.
- Fleurbaey, M., & Voorhoeve, A. (2013). 'Decide as you would with full information' in Eyal, N., Hurst, S. A., Norheim, O. F., and Wikler, D., eds., *Inequalities in health: Concepts, measures, and ethics*, Oxford University Press, 113-128.
- Frick, J. (2013). 'Uncertainty and Justifiability to Each Person' in Eyal, N., Hurst, S. A., Norheim, O. F., and Wikler, D., eds., *Inequalities in health: Concepts, measures, and ethics*, New York: Oxford University Press, 129.
- Frick, J. (2015). 'Contractualism and social risk', *Philosophy & Public Affairs*, 43(3), 175-223.
- Fried, B. H. (2018). 'Facing up to Risk', *Journal of Legal Analysis*, 10.
- Gordon-Solmon, K. (2019). 'Should Contractualists Decompose?', *Philosophy & Public Affairs*, 47(3), 259-287.
- Halstead, J. (2016). 'The numbers always count', *Ethics*, 126(3), 789-802.
- Hare, C. (2013). *The limits of kindness*. Oxford University Press.
- Hirose, I. (2015). *Moral aggregation*. Oxford University Press, USA.
- Horton, J. (2018). 'Always Aggregate', *Philosophy & Public Affairs*, 46(2), 160-174.
- Horton, J. (2017). 'Aggregation, complaints, and risk', *Philosophy & Public Affairs*, 45(1), 54-81.
- John, S. D. (2014). 'Risk, contractualism, and Rose's "prevention paradox"', *Social Theory and Practice*, 40(1), 28-50.
- Johnson, C. M. (2019). 'How Deontologists Can Be Moderate (and Why They Should Be)', *The Journal of Value Inquiry*, 1-17.
- Kamm, F. M. (1998). *Morality, Mortality Volume I: Death and Whom to Save From It*. Oxford University Press.
- Kamm, F. M. (2008). *Intricate Ethics: Rights, Responsibilities, and Permissible Harm*. Oxford University Press.
- Kelleher, J. P. (2014). 'Relevance and non-consequentialist aggregation', *Utilitas*, 26(4), 385-408.
- Kirkpatrick, J. R. (2018). 'Permissibility and the Aggregation of Risks', *Utilitas*, 30(1), 107-119.
- Lange, B. (2017). 'Restricted prioritarianism or competing claims?', *Utilitas*, 29(2), 137-152.
- Lazar, S. (2018). 'Limited Aggregation and Risk', *Philosophy & Public Affairs*, 46(2), 117-159.
- Liao, S. M. (2008). 'Who is afraid of numbers?', *Utilitas*, 20(4), 447-461.
- Nagel, T. (1979). 'Equality' in *Mortal Questions*. Cambridge University Press, 106-127.
- Otsuka, M. (2015). 'Risking life and limb' in Cohen, I. G., Daniels, N., and Eyal, N., eds., *Identified versus statistical lives: an interdisciplinary perspective*, New York: Oxford University Press, 77-93.
- Otsuka, M., & Voorhoeve, A. (2009). 'Why it matters that some are worse off than others: an argument against the priority view', *Philosophy & Public Affairs*, 37(2), 171-199.
- Parfit, D. (1984). *Reasons and persons*. Oxford University Press.
- Parfit, D. (2011). *On what matters: volume one* (Vol. 1). Oxford University Press.
- Rawls, J. (1971) *A Theory of Justice*. Harvard University Press.

- Rowe, T. (2019). 'Risk and the unfairness of some being better off at the expense of others', *Journal of Ethics and Social Philosophy*, 16(1).
- Rowe, T., & Voorhoeve, A. (2018). 'Egalitarianism under severe uncertainty', *Philosophy & Public Affairs*, 46(3), 239-268.
- Ruger, K. (2018). 'On ex ante contractualism', *J. Ethics & Soc. Phil.*, (13)3, 240
- Scanlon, T. (1998). *What we owe to each other*. Harvard University Press.
- Scheffler, S. (2018). *Why Worry about Future Generations?*. Oxford University Press.
- Suikkanen, J. (2019). 'Ex ante and ex post contractualism: A synthesis', *The Journal of Ethics*, 23(1), 77-98.
- Swanton, C. (2003). *Virtue ethics: A pluralistic view*. New York: Oxford University Press.
- Tadros, V. (2019). 'Localized restricted aggregation', in Sobel, D., Vallentyne, P., and Wall, S., eds., *Oxford Studies in Political Philosophy Vol. 5*, New York: Oxford University Press, 17.
- Taurek, J. M. (1977). 'Should the numbers count?', *Philosophy & Public Affairs* 6(4), 293-316.
- Temkin, L. S. (2012). *Rethinking the good: Moral ideals and the nature of practical reasoning*. Oxford University Press.
- Tenenbaum, S. (2017). 'Action, deontology, and risk: against the multiplicative model', *Ethics*, 127(3), 674-707.
- Thompson, C. (2018). 'Rose's prevention paradox', *Journal of Applied Philosophy*, 35(2), 242-256.
- Tomlin, P. (2017). 'On limited aggregation', *Philosophy & Public Affairs*, 45(3), 232-260.
- Van Gils, A. & Tomlin, P. (forthcoming) 'Relevance rides again? Aggregation and local relevance', in Sobel, D., Vallentyne, P., and Wall, S., *Oxford Studies in Political Philosophy*, Vol. 6.
- Voorhoeve, A. (2014). 'How should we aggregate competing claims?', *Ethics*, 125(1), 64-87.
- Voorhoeve, A. (2016). 'Why one should count only claims with which one can sympathize', *Public Health Ethics*, 10(2), 148-156.
- Voorhoeve, A. (2018). 'Balancing small against large burdens', *Behavioural Public Policy*, 2(1), 125-142.
- Walen, A. D. (2019). *The Mechanics of Claims and Permissible Killing in War*. New York: Oxford University Press.
- Walen, A. D. (forthcoming in *Ethics*). Weak Aggregation and Risk. [Received 07.08.2019]
- Wallace, R. J. (2019). *The Moral Nexus* (Vol. 8). Princeton: Princeton University Press.