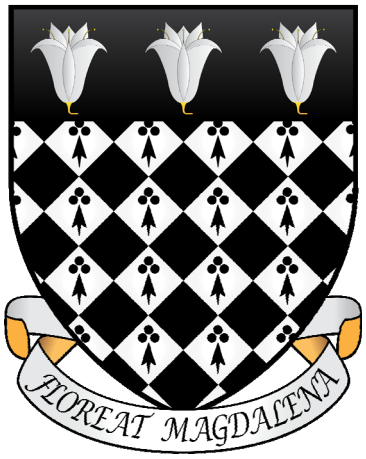
WHAT TO DO WHEN THERE ISN'T ENOUGH

THE FAIR DISTRIBUTION OF SCARCE GOODS





Name: Gerard Vong (OSS ID: 325356)
College: Magdalen
Primary Supervisor: Professor John Broome

Submitted for the Doctor of Philosophy in Philosophy degree.

Word Count (including citations): 54, 805

TABLE OF CONTENTS

WHAT TO DO WHEN THERE ISN'T ENOUGH - THE FAIR DISTRIBUTION OF SCARCE GOODS
ABSTRACT 4

ANTI-MAJORITARIANISM - AGAINST ALWAYS BENEFITING THE GREATER NUMBER 5

SECTION I – QUALITATIVELY EQUAL CONFLICT CASES 6
SECTION II - AGAINST ALWAYS BENEFITING THE GREATER NUMBER 10

WEIGHING UP WEIGHTED LOTTERIES 15

SECTION I – CASES & TERMINOLOGY 18
SECTION II - AGAINST UNWEIGHTED LOTTERIES 20
SECTION III - EVALUATING EXISTING WEIGHTED LOTTERIES 30
III.A – Kamm's Proportional Chances 30
III.B – Timmermann's Individualist Lottery 33
SECTION IV – INTRODUCING COMPOSITION SENSITIVE LOTTERIES 35
IV.A - Composition Sensitive Lottery Variants 37
IV.B - A Composition Sensitive Variant of Timmermann's Individualist Lottery 42
SECTION V - EVALUATING COMPOSITION SENSITIVE LOTTERIES 43
V.A - Convergent Cases 43
V.B - Non Maximal Outcome Groups 48
SECTION VI - THE RELATION OF THE FAIR TO THE RIGHT 52
SECTION VII - DEFENDING LOTTERY PROCEDURES 60
VII.A - Against Pooling 60
VII.B - The 'Grand Lottery of Life' Objection 63
VII.C - Unequal Benefits 68
CONCLUSION 70

CAN'T GET NO SATISFACTION - THE CHANCY SATISFACTION OF MORAL CLAIMS 74

SECTION I - THE ROLE OF LOTTERIES IN FAIRNESS & BROOME'S SYNCHRONIC VIEW 75
I.A - The Role of Chances of Benefiting for Fairness 75
I.B - Broome's Synchronic View of Claim Satisfaction by Lottery 77
I.C - Objections to Broome's Synchronic View 78
SECTION II - ALTERNATIVE VIEWS OF CLAIM SATISFACTION 87
II.A - Diachronic Alternatives 87
II.B - The Dual Structure View 90
II.C - Defending the Dual Structure View 94
SECTION III - PRACTICAL IMPLICATIONS 99

VALUING THE FEW OVER THE MANY 108

SECTION I - DISPARATE CONFLICT CASES & THE SECOND STANDARD VIEW 109
SECTION II - DEFENDING THE SECOND STANDARD VIEW 115
SECTION III - RESOLVING AXIOLOGICAL INCONSISTENCY 122
CONCLUSION 142

MAKING BETTERNESS BEHAVE - CONSTRUCTING A TRANSITIVE DEONTIC ORDERING FROM A NON-TRANSITIVE AXIOLOGICAL ORDERING 145

SECTION I - ARGUMENTS FOR THE NON-TRANSITIVITY OF BETTER THAN 147
SECTION II - SUPPOSED COSTS OF NON-TRANSITIVITY OF BETTER THAN 151
II.A - Money Pumps 151
II.B - The Practical Dilemmas of Never Acting Rightly 152
II.C - Removing the Justification of Pairwise Comparison 154
SECTION III - MITIGATING THE COSTS 155
III.A - Motivating The Difference between the Axiological and the Deontic 155
III.B - The Commitment Deontic Principle & Kamm's Axiological Prioritization 158
III.C - Non-sequential Cases That Resist Prioritization 165
III.D - Constructing a Transitive Deontic Ordering 167

SECTION IV - THE MERITS OF MY TRANSITIVE DEONTIC ORDERINGS 171
IV.A - Theoretical Merits 171
IV.B - Practical Merits of the Ordering 173
IV.C - Solving the Money Pump Problem 174
SECTION V - OBJECTIONS TO MY DEONTIC ORDERING METHODS 184
V.A - Violation of Pareto 184
V.B - Option Individuation 186
V.C - Limitations of Actual Agents 190
V.D - Are my Two Deontic Ordering Methods Ad Hoc? 191
CONCLUSION 195

What to Do When There Isn't Enough - The Fair Distribution of Scarce Goods Abstract

Name: Gerard Vong (Magdalen College)
Submitted: Michaelmas Term 2012
Degree: DPhil in Philosophy

My DPhil submission consists of a series of papers on related topics on the moral philosophy of scarce benefit distribution. It focuses on two types of scarce benefit distribution case.

The first type occurs when which all potential beneficiaries of a good each have an equally strong moral claim on an equal benefit from the resource but scarcity or indivisibility prevents us from benefiting all potential beneficiaries. Call these cases *equal conflict cases.* In 'Anti-Majoritarianism', I argue against the view defended by both utilitarians and non-utilitarians that in equal conflict cases you always ought to give the benefit to as many people as possible. I argue that doing so is neither morally right nor fair. In 'Weighing Up Weighted Lotteries', I argue that the philosophical debate between unweighted and weighted lottery benefit distribution procedures has been misconceived and that fairness requires us to use a new kind of weighted lottery that I call the *exclusive composition-sensitive lottery*. In 'Can't Get No Satisfaction', I defend a new view that I call the *dual-structure view* about how lotteries satisfy potential beneficiaries' claims in equal conflict cases and highlight the implications of that view for the distribution of donor corneas to those who have suffered corneal degeneration.

The second type of this distributional problem occurs when we can either benefit a very large number of potential beneficiaries with a very small benefit (call these *the many*) or a very small number of potential beneficiaries with a very major benefit (call these *the few*). In "Valuing the Few Over the Many" I argue that there are cases where not only ought we to benefit *the few* over *the many* no matter how numerous *the many* are, but it is also better to do so. However, this conclusion can be shown to conflict with a number of widely held tenets of value theory. I evaluate different ways of accommodating these intuitions and argue that in some contexts, benefits are not of finite value. The view I defend in 'Valuing the Few Over the Many', combined with some intuitively plausible axiological claims, is inconsistent with the transitivity of the 'better than' relation. In 'Making Betterness Behave' I argue that for what I call the *conditional non-coextensive* thesis: if 'better than' is not transitive, one ought to take the position that 'more reason to bring about rather than' *is* transitive. I argue that one can generate a transitive 'more reason to bring about rather than' deontic ordering from a non-transitive axiological ordering in a principled way. This deontic ordering avoids the major practical objections (money pumps, moral dilemmas and threats to practical reasoning) to non-transitivity of the 'better than' relation.

Anti-Majoritarianism - Against Always Benefiting the Greater Number

One of the most pressing issues in both ethics and economics is how we ought to distribute scarce, indivisible resources. This paper criticizes the most prominent solution to a particular variant of this problem, namely when all potential beneficiaries of the resource each have an equally deserving claim on an equal benefit from the resource, and argues that we ought to find a fairer approach. Many prominent moral and economic theories, both utilitarian and non-utilitarian, mistakenly claim that in such cases, you always ought to give the benefit to as many people as possible. Call such theories majoritarian theories. I will argue that acting according to majoritarian theories does not appropriately satisfy people's claims and for this reason is both morally objectionable and unfair. To this end, the paper is structured as follows. Section I will introduce some philosophically useful terminology and specify the kinds of distribution case this paper will focus on. Section II will argue against majoritarian theories by showing that such theories lead to objectionable discrimination in iterative cases and that because they focus solely on the number of actual benefits distributed, they forego an important aspect of fairness. I turn now to elucidate the types of cases this paper focuses on.

Section I – Qualitatively Equal Conflict Cases

Consider the following case:

> *1v5 Lifeboat*: One person, Alice, is stranded on one rock R1 and five other people (Betty, Celine, Diane, Edna & Frances) are stranded on another rock R2. You have a seaworthy lifeboat with which you can easily save the occupants of one rock before the tide rises and drowns everyone you fail to save. All six potential people stranded on the rocks are young and innocent and will lose a great deal by dying. In this case and others I will discuss in this paper, no one knows anyone else and no one differs in any other morally relevant way.

Call this case *1v5 Lifeboat,* where the 1 and 5 are indicative of the number of people on rocks R1 and R2 respectively. Call the different groups that can be benefited (in this case, the occupants on R1 or R2), outcome groups. Given the opportunity to save either one of the outcome groups, our intuitions are that we ought to save the five on rock R2.

1v5 Lifeboat is important here for two reasons. Firstly, it is a particular version of a general problem, namely how to distribute benefits when all potential

beneficiaries have a qualitatively equal moral claim to an equal benefit[1] that cannot be distributed amongst all equally worthy claimants. Call this an *equal conflict case*.

The reasons I have eschewed discussion of resources in favour of benefits are two fold. Benefits are more general than resources, as resources have a number of connotations I want to avoid. For example, resources are often viewed as tangible goods owned by a particular person or organization. More importantly, some resources, such as military resources, can often be used to harm people, and in this paper I want to focus on the ethics of benefit distribution, not harm distribution.

These are *conflict* cases because they are cases in which choosing to benefit some set of claimants prevents one from benefiting other claimants. By claim, I mean a particular kind of reason, namely acts that are owed to a particular individual or group. Claims are patient-relative reasons, where a patient is a person or thing that undergoes an action of another (this meaning of 'patient' contrasts with 'agent'). To illustrate the distinction between reasons that are claims and those that are not, note that there are reasons to benefit a person that are not acts owed to that person and there are reasons to benefit a person that are acts owed to that person. For example, if you promise X to benefit Y, you have a reason to benefit Y, but this reason is owed to X and not Y.

[1] In *1v5 Lifeboat*, the benefit is preventing their premature death.

In this paper, I intend to remain largely neutral about which reasons are claims and how claims are established, maintained, strengthened or weakened. My paper, does, however, draw upon a number of thought experiments that are intended to be examples of the *equal conflict cases* and as such it makes some assumptions about claims. For example, in the life rescue cases I discuss, I assume that all of the potential beneficiaries each have an equal claim to be rescued. One plausible way in which such a claim could arise is through the right to life creating a (defeasible) duty to save each person. When it is easy to save any individual, as in *1v5 Lifeboat*, it is plausible that each individual has a claim to be saved. Note also that the potential beneficiaries have the negative right not to be treated unfairly and this gives them a claim not be treated unfairly. What is important in the cases I discuss are the different possible ways to distribute benefits and the existence of qualitatively equal claims and nothing else rests on the particular narrative details of each case. Readers with differing views on how claims arise can construct structurally similar variants of my cases in which all beneficiaries have equal claims on the benefit.

How to decide whom to benefit in *equal conflict cases* is a prevalent moral problem, evident not just in rescue cases, but also in medical resource distribution, financial obligations and numerous other areas of public policy and private moral behaviour. For example, how should we distribute donor kidneys from cadaveric sources when demand for medical operations far outstrips supply? What should you do if you promised two equally good friends that you would drive them to Wimbledon tomorrow but find out just before leaving that

you only have car space for one and cannot compensate whoever you leave behind? Often such cases arise from insufficient resources to benefit all claimants. However, sometimes such a problem arises due to urgency, claimant accessibility[2] or indivisibility of a plentiful benefit[3].

The second reason *1v5 Lifeboat* is important here is that our strongly held intuitions in this case show that numbers do matter – all other things being equal, benefiting more claimants in an equal claims conflict case seems better than benefiting fewer claimants. It is important for moral theory to account for this strongly and widely held intuition. Standard act consequentialist theories (e.g. act utilitarianism) initially seem to get the right answer by aggregating the value of the five individuals on R2. According to such theories, saving the greater number is the right action because the value of saving the five is greater than just the value of saving a single life. Nonconsequentialist procedures, such as the Scanlon[4]-Kamm[5] 'tie-breaking' procedure, also recommend benefiting the greater number in equal claims conflict cases. While theories that always recommend benefiting the greater number in equal claims conflict cases (call such theories majoritarian theories) capture our intuitions in simple cases such as *1v5 Lifeboat*, I will argue in the next section that all majoritarian theories, whether consequentialist or nonconsequentialist, have two fundamental

[2] For example, claimants are in different places so that we cannot send the benefit to all claimants.
[3] For example, you owe five people lunch, but you can only fulfil this obligation with one claimant and even then, only by giving them a voucher that entitles them and only them free meals for a year.
[4] Thomas Scanlon, *What we owe to each other,* (Cambridge, Mass. ; London: Belknap Press of Harvard University Press, 1998), 229-234.
[5] Frances Myra Kamm, *Morality, Mortality: Volume 1,* (New York ; Oxford: Oxford University Press, 1993), 116-119.

failings which make them unable to account for our strongly and widely held intuitions in a number of more complex equal claims conflict cases that I will discuss.

Section II - Against Always Benefiting the Greater Number

Always acting according to the recommendations of majoritarian theories can lead to long term discrimination against minorities and for this reason is objectionable. This is evident when the same groups of claimants are repeatedly involved in equal claims conflict cases. To see this, imagine we have to choose whether to benefit group G_1 or group G_2, where the individuals in both groups have qualitatively equal claims on the benefit but we can only choose to benefit one of the groups. We can give out this benefit once every month for a year. Say G_1 has 15 members and G_2 has 10 members – if we always choose to benefit the larger group, then G_1 will benefit consistently over the year, whereas G_2 will never receive that benefit. This is both unfair and objectionable and gives us reason to doubt the plausibility of majoritarian theories. However, as I will argue, it is not just in iterative equal conflict cases that majoritarian theories affront our intuitions – even in individual cases in which the exact same claimants will *never* be in an equal conflict case again, we are hesitant to straightforwardly benefit the greater number of claimants.[6]

[6] Scanlon's tie-breaking procedure is not designed for equal conflict cases in which the same claimants are repeatedly in conflict cases which other. See Scanlon, *What we owe to each other*, 397n38.

Furthermore, by focusing on the actual benefits distributed in equal conflict cases, majoritarians forego an important feature of fairness by not giving any but the largest outcome group(s) any chance of benefiting. This is because fairness does not only depend on the actual benefits distributed but also on the chances of receiving the benefit. To see this, compare the following two courses of action in which the actual benefit distribution is identical but beneficiaries' ex ante chances of receiving the benefit are different. While this case involves equally strong claims, this is a non-conflict case because it is possible to give all potential beneficiaries the benefit. There are two courses of action, the first of which, B_1, involves giving all claimants a 100% chance of receiving the benefit and all claimants receive the same benefit. In the other, B_2, claimants are given different vastly different chances of receiving the benefit – some claimants are given 1% chance of benefiting, others 60%, others 100% (the precise percentages do not matter, the variation does). However, just by pure luck, all of the claimants receive the benefit, even those with a very low chance of receiving the benefit. B_2 is intuitively less fair than B_1, even though in both all potential beneficiaries actually receive the benefit. This illustrates that fairness depends not just on benefits distributed but also on chances of receiving the benefit. To see how this majoritarian failing manifests itself in equal claims conflict cases, consider,

10,000v10,001 Cruise Ship. Imagine a rescue team is able to save the occupants of one of two giant cruise ships before the other cruise ship sinks and its occupants die in icy, Antarctic waters. One cruise ship (S_1) has ten

thousand occupants, the other (S_2) has ten thousand and one occupants. What should the rescue team do?

In this equal claims conflict case, it again does not seem straightforward that one should just save the greater number and give no chance for the 10,000 on S_1 to be saved.[7] Not only do majoritarian theories suggest that we ought to save S_2, condemning the 10,000 on S_1 to death with no chance of survival, they objectionably appear to regard such an action as *obviously* the right answer.[8] I will provide another type of case which illustrates the aforementioned majoritarian failing in which our intuitions against straightforwardly saving the greater number will motivate looking for a non-majoritarian theory. Consider,

8v7v7v7v7v7v6v6v5v5v4v4v4v4v4v4v4v3v3v1 Archipelago. Imagine a rescue team can save the inhabitants of one island in a thinly populated, typhoon-threatened archipelago before the typhoon kills all remaining island inhabitants.[9] Stipulate that there are 100 people scattered over 20 islands as illustrated in the case's name. One island has eight inhabitants, which is more inhabitants than any other. What should the rescue team do?

[7] This case is similar to a case presented by Kamm which she attributes to Peter Unger. Kamm, *Morality, Mortality: Volume 1*, 103, 121n4.

[8] This echoes a criticism made of utilitarianism by Bernard Williams throughout his contribution to John Smart and Bernard Williams, *Utilitarianism; for and Against*, Cambridge University Press, 1973)

[9] This example is based on a case from David Wasserman and Alan Strudler, "Can a Nonconsequentialist Count Lives?" *Philosophy and Public Affairs* 31, no. 1 (2003): 71-94.

Any majoritarian theory is intuitively wrong and unfair because it would recommend the abandonment of the *majority* of the islanders by recommending only to rescue the island of eight as obviously the right thing to do. That is, 92 of the 100 people are not given any chance of survival on such an outcome group selection procedure. To see the importance of this, consider a majoritarian's act prescription to save the island of 908 in *908v7v7v7v7v7v6v6v5v5v4v4v4v4v4v4v4v3v3v1 Archipelago*. While the number of individuals (92) given no chance of survival is the same as the first *Archipelago* case, it is intuitively less unfair as those given no chance of survival are a minority of all claimants. The relative size of outcome groups is important. As such, an intuitively plausible and fair theory of benefit distribution in equal claims conflict cases must account for both individuals and the relative size of groups.

Each of these cases (G_1&G_2, *10,000v10,001 Cruise Ship* & *8v7v7v7v7v7v6v6v5v5v4v4v4v4v4v4v4v3v3v1 Archipelago*) are just particular examples of more general types of equal claims conflict cases in which our intuitions clash with majoritarian theory's act prescriptions. We are hesitant to always benefit the greater number in iterative equal claims conflict cases or where doing so prevents us from giving any benefit to a slightly smaller group of claimants, especially when there are particularly many claimants in the largest outcome groups (as in *10,000v10,001 Cruise Ship*) or when the largest possible group that we can benefit is a minority of all worthy claimants (as in

8v7v7v7v7v6v6v5v5v4v4v4v4v4v4v3v3v1 Archipelago). These generalizable cases are indicative of a fundamental flaw in majoritarian theory, namely by focusing solely on the actual benefits distributed and always benefiting the greater number, they forego giving claimants who are not in the largest outcome group any chance of benefiting. This is intuitively unfair and as a result, majoritarians appear to be missing something of moral relevance in their moral theory. The arguments of this paper provide good reasons to look for a non-majoritarian theory that better respects fairness. I take up that task in my paper, 'Weighing up Weighted Lotteries'.

Bibliography

Kamm, Frances Myra. *Morality, Mortality: Volume 1.* New York ; Oxford: Oxford University Press, 1993.

Scanlon, Thomas. *What we Owe to each Other.* Cambridge, Mass. ; London: Belknap Press of Harvard University Press, 1998.

Smart, John, and Bernard Williams. *Utilitarianism; for and Against.* Cambridge University Press, 1973.

Wasserman, David, and Alan Strudler. "Can a Nonconsequentialist Count Lives?" *Philosophy and Public Affairs* 31, no. 1 (2003): 71-94.

Weighing Up Weighted Lotteries[10]

One of the most pressing issues in both ethics and economics is the fair distribution of scarce, indivisible resources. This paper discusses a particularly difficult instance of this problem, namely when all potential beneficiaries of the resource each have an equally strong claim on an equal benefit from the resource[11].

This instance is not merely of theoretical interest, but is relevant to many practical questions. How should we fairly distribute donor kidneys from cadaveric sources when demand for medical transplants far outstrips supply? What is the fair distribution of measles vaccinations that have been allocated to a developing country when there are insufficient resources to vaccinate everyone in that country? What is the fair allocation of scarce intensive care beds during a flu pandemic?

In such cases, I argue that a new, fairer approach to distribution is needed. Given that we cannot give all claimants the benefit they deserve, it is fair to give claimants appropriate chances at receiving the benefit as partial surrogates for

[10] I have benefited greatly through discussion and presentation of parts of this material. I am thankful to Stuart Armstrong, Katharina Berndt, Dan Brock, John Broome, Krister Bykvist, Richard Chappell, Roger Crisp, Tom Dougherty, Daniel Halliday, Matthew Liao, David Miller, Michael Otsuka, Derek Parfit, Andrew Schroeder, Peter Stone, Noel Swanson, Larry Temkin, Alex Voorhoeve, Dan Wikler and Dominic Wilkinson as well as audiences at Harvard University, Hokkaido University, Princeton University, Oxford University, the Institute of Philosophy at the University of London and the University of Sydney.

[11] I use resource here and throughout the paper in a broad sense to include anything that can provide a benefit.

actually receiving the benefit[12]. The use of lotteries to distribute scarce benefits has a long history from Ancient Greece to modern medical distribution[13], legal cases[14] and ticketing for popular events such as the Olympics. The current philosophical debate on using lotteries to distribute benefits is between unweighted and weighted lotteries. Unweighted lotteries always give all claimants an equal chance of benefiting, whereas weighted lotteries sometimes give claimants an unequal chance of benefiting[15]. The received view is that in cases of scarcity such as these, fairness requires us to use unweighted lotteries to distribute benefits[16]. I will argue that the current debate between unweighted and weighted lottery theorists is misconceived as there is an important range of cases in which one *cannot* give potential beneficiaries equal chances of benefiting. Therefore a lottery theorist that wants to account for all cases *must* adopt a weighted lottery. Furthermore, none of the currently proposed weighted lotteries account for these more complex cases. Instead, fairness requires us to run a new kind of weighted lottery that I will call the *composition sensitive lottery* because it is sensitive to the composition of groups of potential beneficiaries in a way that other lotteries are not.

---

[12] Here, I adopt John Broome's view of why giving claimants appropriate chances of benefiting are a fair surrogate for actually benefiting them, though I will later argue against his view that equal chances of benefiting are the most fair. Broome, John "Fairness", in John Broome, *Ethics Out of Economics,* (Cambridge: Cambridge University Press, 1999), 111-122.

[13] Barbara Goodwin, *Justice by lottery,* 2nded. (Exeter: Imprint Academic, 2005), 53-54.

[14] Leo Katz, *Bad acts and guilty minds : conundrums of the criminal law,* (Chicago ; London: University of Chicago Press, 1987)

[15] This is the case in one off distribution cases. When an unweighted lottery is used to distribute benefits in a number of iterated benefit distribution cases where the potential beneficiaries change, it is possible for it to give some potential beneficiaries a greater chance of benefiting.

[16] Proponents, for example, include John Broome, Iwao Hirose and Michael Otsuka. See Broome, John "Fairness", in Broome, *Ethics Out of Economics,* 111-122.; Michael Otsuka, "The Fairness of Equal Chances," (Unpublished Manuscript); Iwao Hirose, "Weighted Lotteries in Life and Death Cases," *Ratio* 20, no. 1 (2007): 45-56.

The paper is structured as follows. Section I will introduce some useful terminology, explicate the notion of fairness most relevant to this paper and specify the kinds of distribution case this paper will focus on. Section II will argue against unweighted lotteries by showing that they have implausible and unfair implications because by explicitly focusing only on equal chances to individual claimants, they have ignored the importance of groups and group composition. Furthermore, I argue that there are a heretofore unconsidered class of cases that unweighted lotteries cannot account for and any attempt to account for them with a lottery *must* be a weighted lottery account. Section III will outline and object to all existing weighted lotteries in the literature. Section IV will provide a taxonomy of new weighted lottery approaches. This new taxonomy arises when we realise that the relevant philosophical debate regarding the application of lotteries to these cases ought not to be between unweighted lotteries and weighted lotteries, but instead between different lotteries that can overcome the objection I raise in Sections II and III. Section V evaluates the new lotteries presented in Section IV and advocates one lottery as the lottery that fairness requires. Section VI discusses the relation between what is fair and what is right in these cases, arguing that what is fair is not coextensive with what is right. Section VII defends lottery distribution procedures from three important objections. I turn now to elucidate the types of cases this paper focuses on.

Section I – Cases & Terminology

Consider,

> *2v3 Lifeboat*: two people Alice & Betty (A&B) are stranded on one rock and three other people Celine & Diane & Edna (C&D&E) are stranded on another rock. You have a seaworthy lifeboat with which you can easily save the occupants of one rock before the tide rises and drowns everyone you fail to save. All of A through E are young and innocent and will lose a great deal by dying. No one knows anyone else and no one differs in any other morally relevant way.

2v3 Lifeboat is a particular version of a general problem I want to discuss, namely, how to distribute benefits when all potential beneficiaries have a qualitatively equal moral claim to an equal benefit[17] that cannot be distributed amongst all equally worthy claimants. Call this the problem of benefit distribution in *equal conflict cases*. It is a *conflict* case because it is one in which choosing to benefit some set of claimants prevents one from benefiting other claimants. By claim, I mean a particular kind of reason, namely an act or good that is owed to a particular individual or group. Claims are patient-relative reasons, where a patient is a person or thing that undergoes an action of

[17] In *2v3 Lifeboat*, the benefit is preventing their premature death.

another (this meaning of 'patient' contrasts with the meaning of 'agent')[18]. To illustrate the distinction between reasons that are claims and those that are not, note that there are reasons to benefit a person that are not owed to that person and there are reasons to benefit a person that are owed to that person. For example, if you promise X to benefit Y, you have a reason to benefit Y, but this reason is owed to X and not Y. Call the different groups of individuals that can be benefited outcome groups. The numbers 2 and 3 in *2v3 Lifeboat* are indicative of the outcome group sizes. In *2v3 Lifeboat* A&B is one outcome group and C&D&E are another outcome group.

In this paper, the fairness I refer to is *procedural* fairness. In particular, it is the appropriate procedure by which people's claims should be satisfied. Remember here that by definition, claims are only one particular kind of reason. As a result, what is procedurally fair is not necessarily all reasons considered morally right[19].

Claimants' *ex ante* objective chances of benefiting are one important feature relevant to procedural fairness. To see this, consider,

> *The Chancy Doctor*: A doctor can conduct lifesaving surgery on two of her equally badly off patients in two separate operations, who are the only two

[18] Sometimes patients have claims on no particular agent, for example when a claimant could have their premature death easily prevented by any one of a number of potential agents.

[19] As Brad Hooker notes in his paper, 'Fairness', there is certainly one broad usage of the word 'fair' that corresponds to all things considered moral judgments. I agree with Hooker that this broad use of the term fair is of limited use in moral philosophy as we already have terms for all things considered moral judgments. Brad Hooker, "Fairness," *Ethical theory and moral practice* 8, no. 4 (2005): 329-352.

patients she can currently help, which she can do easily with guaranteed success. Instead of just conducting the surgery on both patients, she decides to roll two separate, genuinely random, six sided dice. If and only if the first die rolls a 1, she will choose to conduct surgery on the first patient. If and only if the second die rolls between 1-5, she will choose to conduct surgery on the second patient. The first die rolls 1 and the second 4, so she actually conducts successful surgeries on both patients.

Even though the *Chancy Doctor* gives the same lifesaving benefit she owes to both patients, her action is patently unfair because of the differences in the *ex ante* chances of benefiting she gives to the two patients. It would be even less fair if she gave one patient a 100% chance to be treated but gave no chance to the other. Because she can save both of their lives easily with guaranteed success, fairness requires her to give both patients an equal 100% chance of benefiting. However, the fair chances of benefiting in equal conflict cases are far more complex - to this question I now turn.

Section II - Against Unweighted Lotteries

An unweighted lottery would give everyone an equal chance of being saved in *2v3 Lifeboat* by giving 50% chance of saving A & B and a 50% chance of saving C, D & E. There is something initially appealing about giving every equally worthy claimant an equal chance of benefiting and many philosophers advocate equal

chances as the fair method of distribution in equal conflict cases including Broome[20], Hirose[21], Otsuka[22] and Peterson[23]. Some philosophers have even *defined* fair lotteries in terms of equiprobable outcomes[24]. Strikingly, Taurek argues that running an unweighted lottery in equal conflict cases is the morally right procedure[25].

However I will argue that the initial appeal of unweighted lotteries is the same kind of defeasible initial appeal evident in simple additive versions of act hedonic consequentialism - the devil is in the details. One problem in the lottery debate is that unweighted lottery theorists fail to properly consider more difficult, more complex equal conflict cases. For example, Lang writes "I will typically focus on coin-tossing here, on the harmless assumption that we are considering cases in which only two groups of individuals are involved. Of course, some other random selection device must be used where there are more than two rescue groups to be considered, but Taurek's argument [for an unweighted lottery] can be comfortably adapted to cover these conflict cases"[26]. Not harmless - not only can unweighted lotteries not be comfortably adapted to some more complex conflict cases, in some cases I will discuss, unweighted lotteries are *impossible* to run and weighted lotteries are the only lottery option available. While unweighted lotteries offer equality of chances to

[20] Broome, John "Fairness", in Broome, *Ethics Out of Economics,* 111-122.
[21] Hirose, *Weighted Lotteries in Life and Death Cases,* 45-56.
[22] Otsuka, *The Fairness of Equal Chances,*
[23] Martin Peterson, "The Mixed Solution to the Number Problem," *Journal of Moral Philosophy* 6, no. 2 (2009): 166-177.
[24] Peter Stone, "On Fair Lotteries," *Social Theory and Practice* 34, no. 4 (2008): 573-590.
[25] John Taurek, "Should the Numbers Count?" *Philosophy and Public Affairs* 6, no. 4 (1977): 293-316.
[26] Gerald Lang, "Fairness in life and Death Cases," *Erkenntnis* 62, no. 3 (2005): 321-351.

receive the benefit, I will show that they are both implausible and unfair. By focusing only on individual chances of benefiting, such lotteries give implausible prescriptions about what is most fair when the possible benefit distributions become more complex than simple mutually exclusive outcome groups of roughly equal size.

One way in which unweighted lotteries are unfair is evident when you consider that the outcome group selection procedure it recommends is sensitive to the addition of new outcome groups but insensitive to additional people's claims when those people join existing outcome groups[27]. To see this, consider,

> *33v33 Train Cars*: two train cars containing 33 people each are hurtling towards an incomplete bridge over a deep chasm and a rescue team can save the occupants of one car before the entire train falls down the chasm and all of its remaining occupants are killed. The occupants cannot escape from the train by themselves.

[27] It should be noted that it is the *outcome group selection procedure* that is insensitive to additional people's claims when those people join existing outcome groups. That is, the method of selecting which of the two outcome groups to save is insensitive to the actual numbers in the train cars (or lifeboats) - it would not change no matter what non-zero number of people were in the outcome groups. Scanlon and Kamm make the *similar* claim that additional people make no difference at all to someone acting according to Taurek. Otsuka, however, has shown that this is incorrect, as Taurek would advocate saving everyone in the boat that his coin flipping deliberative procedure selects (that is, if the coin flip suggests we should save the individuals on the second rock, in *1v3 Lifeboat* Taurek will save 3 individuals but in *1v4* Taurek will save 4 individuals). My version of the objection avoids this criticism. See Thomas Scanlon, *What we owe to each other,* (Cambridge, Mass. ; London: Belknap Press of Harvard University Press, 1998), 232-233.Frances Myra Kamm, *Intricate ethics : rights, responsibilities and permissible harm,* (New York, N.Y. ; Oxford: Oxford University Press, 2007), 33. Michael Otsuka, "Saving Lives, Moral Theory, and the Claims of Individuals," *Philosophy & Public Affairs* 34, no. 2 (2006): 109-135.

In such a equal conflict case, Broome would recommend flipping an unbiased coin as a fair way to determine which outcome group to save. If the situation were identical except for the addition of a third train car with 33 occupants (thereby becoming *33v33v33 Train Cars*), Broome would allocate each outcome group a 1/3 chance of survival. So far, so good – Broome's outcome group selection procedure is sensitive to the addition of new outcome groups of equally worthy claimants. It would be clearly unfair if the outcome group selection procedure was insensitive to the addition of another equally worthy outcome group – in this case, such a procedure would advocate flipping a coin to select either of the first two train cars to save and give no chance to rescue the 33 in the third car. Broome's outcome group selection procedure, however, is insensitive to the addition of new equally worthy claimants to existing groups, as we can see if we add 33 people to the second train car in *33v33 Train Cars*, changing it to *33v66 Train Cars*. In such a case, it would still recommend the same outcome group selection procedure as *33v33 Train Cars*, namely to flip a coin to determine whether to select the first or second train car. In fact, the outcome group selection procedure is insensitive to any number of positive claimants in each outcome group, as it would recommend the same outcome group selection procedure in *33v33 Train Cars, 1v1 Train Cars* or *1v100 Train Cars*.

Stated more precisely, unweighted lotteries violate the following intuitively supported principle,

Sensitivity: fairness in equal conflict cases requires the outcome group selection procedure to be appropriately changed if claimants are added or removed and thereby the relative sizes of the outcome groups change[28].

Why ought we to believe in *Sensitivity*? *Sensitivity* accounts for the intuitively supported and important moral difference between *1v1 Train Cars* and *1v100 Train Cars.* Each extra individual in the more populated train car could complain that it is unfair that their presence made no difference at all to the outcome group selection procedure. My intuitive judgment would be to agree with them. In *1v100 Train Cars* it would be unfair to the second outcome group containing 100 claimants to give them the same chance of benefiting as the first outcome group containing 1 claimant[29]. The importance of outcome groups and their composition will be further elucidated in my discussion of other more complex equal conflict cases in this section and the next.

I claimed that it was a virtue of Broome's view of giving equal chances to receive the benefit that it is sensitive to the addition of new outcome groups, as was evident in *33v33v33 Train Cars*. While their procedure is initially plausible when

[28] The second conjunctive clause is important so that the outcome group selection procedure does not change when equal numbers of new claimants are added to all outcome groups. To see this, consider *5v5 Lifeboat* compared to *6v6 Lifeboat*. Furthermore, note that this principle is open to multiple interpretations of what the appropriate change is. Unweighted lotteries offer no change at all in such cases. I will develop a view on which change is appropriately fair in Sections IV and V.

[29] Furthermore, consider the difference between *10,000v10,001 Lifeboat* and *2v3 Lifeboat*. In both cases, the difference between the two outcome groups is only one person's life, but the relative size of the two groups is much closer in the former case. Intuitively, what it is fair to do in these cases is different and *Sensitivity* can be used to explain this. A theory that only considers the absolute, but not relative, difference in the size of the outcome groups cannot account for this change.

we consider *33v33v33 Train Cars*, it is sensitive to the addition of new outcome groups in an unfair way. This can be seen by considering a situation that begins like *33v33 Train Cars*, except there is an additional train car which is empty (strictly speaking it is *33v33v0 Train Cars*). Now imagine that one lone claimant, due to the jostling of the cars, is thrown into the third train car without being hurt at all but is unable to rejoin their previous outcome group before the train falls off the chasm. The equal conflict case is now *33v32v1 Train Cars* and the creation of a (numerically positive) third outcome group results in giving all individuals a one third chance of being saved. But why should the chances of survival of over 98% of all claimants (the groups of 32 and 33) drop by a third of their previous chance[30] to approximately 33.33% just because one person has moved in a way that prevents them from being benefited with others? This excessive change in chances to survive is intuitively unfair. This is especially unfair to the group of 33 claimants, who unlike the group that is now 32 claimants, have suffered such a drastic reduction in their chance of survival even though neither their qualitative nor quantitative claims as a group have decreased and there are no additional claimants (just an existing claimant moving to a new outcome group), so that group's proportion of all claimants remains the same.

Not only is Broome's procedure implausible for the above reasons, it does not tell us what is most fair in a significant class of equal conflict cases because it only tells us to give all claimants an equal chance of receiving the benefit.

[30] That is, a third of 50% is approximately 17%, hence the reduction from 50% to 33% is approximately a drop of a third.

Imagine we are in a equal conflict case where we *have to* choose to benefit only one of the following groups of claimants: A&B, A&C, A&D or E. This is a particular equal conflict case where, because it is possible to benefit some claimants (in this case, A) in more than one outcome group, we cannot give all claimants equal chances of benefiting. Call these *unequal chance conflict cases*. A's chance of benefiting is the sum of the chances that B and C and D have of benefiting and therefore it is impossible to give all claimants an equal chance of benefiting. In such cases, Broome's view does not tell us which benefit distribution is most fair.

It would be implausible for Broome to say that no benefit distribution is more fair in such unequal chance conflict cases, for clearly some actions are worse and less fair than others. It would, for example, be both intuitively unfair and wrong to give E a 100% chance of being benefited. This indeterminacy in unequal chance conflict cases actually forces proponents of unweighted lotteries to adopt a weighted lottery if they are going to allocate any chances at all. This is because in these cases it is impossible to give all claimants an equal chance of benefiting. The weighted lottery that would be most similar to an unweighted lottery is to claim that the most fair procedure to adopt in equal conflict cases is instead to provide as close as possible to equal[31] chances to receive the benefit. While this would provide a view that is informative in unequal chance conflict cases, it will lead to implausible implications in other cases that I will discuss next.

[31] Perhaps mathematically measured as lowest standard deviation or utilising economic measures of inequality, such as the Gini Coefficient.

Not only is the equal chances procedure indeterminate in unequal chance conflict cases, it provides an implausible and intuitively unfair answer to the problem of benefit distribution in some equal conflict cases. Imagine an equal conflict case where we can choose to save the lives of only one of the following groups of claimants: A&B, A&C, A&D, E or no claimants. Again, this is a particular version of a general case and A's chance of benefiting is the sum of the chances that B and C and D have of receiving the benefit. The important difference with this case and *unequal chance conflict cases* is that because we can choose to benefit no claimants, it is possible to give all claimants an equal chance of receiving the benefit by letting everyone die (i.e. giving everyone 0% chance of benefiting). Call such cases *equal zero benefit cases*. In this case, Broome's view would claim that letting everyone die is the fair thing to do. In print, Broome has already conceded that it is equally fair on his view whether everyone gets 0% chance of benefiting or 50% chance of benefiting[32]. That equal fairness claim is implausible enough and has been persuasively criticized by Hooker[33]. However what is distinct about my objection is that the case I present is one in which the *only* way of giving people equal chances of benefiting is letting everyone die. Broome has to concede not only that implausible equal fairness claim but also that it is fairer to let everyone die than to give some people a chance at living and actually save some lives. This objection provides a fourth reason why we should reject unweighted lotteries.

[32] Broome, "Fairness" in Broome, *Ethics Out of Economics*, 117-18.
[33] Hooker, *Fairness*, 340-41..

The reason why Broome claims that it is equally fair for all people to have a 0% chance of benefiting in these cases is because he adopts a narrow, *comparative* notion of fairness. For Broome, as long as all equally deserving claimants receive the same treatment, it is perfectly fair. This is an important component of fairness, as it is unfair if equally deserving claimants are treated unequally. On Broome's comparative view, if saints all have equally strong claims on an equally good, flourishing life all suffer equally terribly, it is perfectly fair. It is just as comparatively fair for all of those saints to having an equally excellent flourishing life. However, it is not the only morally relevant notion of fairness. There is an *absolute* notion of fairness, where fairness requires people to get what they are due[34]. This absolute notion can differentiate between a world where all the saints suffer equally and another world where the saints all have an excellent flourishing life. The latter world is more fair than the former world, because while comparatively everyone is as well off as each other within each world, in the latter world the saints get what they are due and in the former world the saints do not. Similarly, this absolute notion of fairness explains why it is intuitively fairest for both patients to get the lifesaving treatment in the *Chancy Doctor case* and unfair to let everyone die in *equal zero benefit* cases. People are due the life saving benefit they have a claim on. In this paper, I aim to find the appropriate way to satisfy potential beneficiaries' claims (i.e. give them what they are due) and argue for a procedure that is all things considered

[34] For a recent treatment of the comparative/absolute distinction, see Larry Temkin's 'Justice, Equality, Fairness, Desert, Rights, Free Will, Responsibility, and Luck' in Carl Knight and Zofia Stemplowska, *Responsibility and distributive justice,* Oxford University Press, 2011)

fair and not merely comparatively fair. This notion takes into account both comparative and absolute fairness.

While I agree with Broome's claim that in equal conflict cases fairness requires we give all claimants some kind of chance of receiving the benefit instead of actually receiving the benefit, I disagree with Broome's claim that fairness requires us always to give all claimants equal chances of being benefited for the reasons I have presented so far in Section II. Broome does not provide a plausible account of what is fair in equal conflict cases because his outcome group selection procedure is insensitive to the addition of equally worthy claimants to existing outcome groups (*1v1 Lifeboat* compared to *1v100 Lifeboat*), is sensitive in the wrong way to the addition of some new outcome groups (e.g. *33v32v1 Train Cars*), is indeterminate in *unequal chance conflict cases* and has implausible implications in *equal zero benefit cases.* Fundamentally, all of these problems arise because by explicitly focusing only on equal chances to individuals, equal chance theorists have ignored fairness to groups by being insensitive to outcome group composition. Fairness is not merely comparative, nor does even the comparative component of fairness require all individuals in equal conflict cases to receive equal chances of benefiting. Given these problems with unweighted lotteries, I now turn to discuss and evaluate weighted lottery alternatives.

Section III - Evaluating Existing Weighted Lotteries

III.A – Kamm's Proportional Chances

According to Kamm's *proportional chances* procedure, which is the most prominent weighted lottery approach[35], each claimant in a equal conflict case is given an equal *baseline chance* of 1 divided by the total number of equally worthy claimants[36] and groups consisting of different claimants can choose to 'pool' (combine) their baseline chances if it is possible to distribute benefits amongst the claimants who consent to be pooled. That is, in *2v3 Lifeboat*, A&B can pool their baseline chances because it is possible to save both of them but A&C or B&D cannot pool their chances as those groups cannot be benefited. A lottery that is weighted according to claimants' pooling choices is then conducted. Thus if A&B pool their chances, they have 40% (20%+20%) chance of being saved and if C&D&E pool their chances, they have a 60% (20%+20%+20%) chance of being saved. To support this pooling claim, Kamm asks us to consider a monetary lottery in which there are six possible winners (A,B,C,D,E,F), but five of them (B-F) are able to share the prize so that when shared it will satisfy each individual's needs just as much as them winning the prize alone[37]. In this

[35] Other writers often refer to a weighted lottery that is practically identical to Kamm's without attributing Kamm. Frances Myra Kamm, *Morality, Mortality: Volume 1,* (New York ; Oxford: Oxford University Press, 1993), 123-143.

[36] The baseline chance is *not* the claimant's chance of receiving a benefit, though it is essential in calculating what any claimant's chance of receiving the benefit will be, as we will see through my discussion of pooling. It is 1 divided by the total number of equally worthy claimants because this gives all claimants the same initial baseline chance, thereby giving each claimant an equal contribution to the outcome group selection procedure. It is analogous to giving a lottery ticket to each claimant.

[37] For example, if the money is jointly won and owned, it can be invested and produce a 500% profit that it is impossible to share with A.

situation, she thinks each individual intuitively has "the right to use it [their baseline chance] as he saw fit"[38].

Kamm's *proportional chances* procedure is fairer than Broome's because it gives each equally worthy claimant an equal impact on the outcome group selection procedure and is thus *sensitive* to the moral differences between *1v2 Lifeboat* and *1v100 Lifeboat*. Comparative fairness in equal conflict cases requires some type of equality because all claimants have equal claim on the benefit, but this equality is not equality of individual's chances of benefiting as I argued in Section II. Kamm's procedure is comparatively fair in equal conflict cases because it accords each individual's claim with equal respect in a particular way, namely by ensuring that all equally worthy claimants make an equal impact on the outcome group selection procedure. This equal impact is ensured by giving all claimants an equal baseline chance. Furthermore, it provides a more intuitively supported and fair solution to cases like *33v32v1 Train Cars* because it gives each and every claimant an equal baseline chance while allocating chances of benefiting that are proportional to outcome group sizes. There are, however, significant problems with Kamm's proposal.

One problem with Kamm's *proportional chances* is that it crucially involves claimants' choice and consent in how the chances to receive the benefit are to be pooled[39]. That is, whether or not baseline chances can be pooled depends on claimants' consent. This leads to problems for the procedure if the claimants

[38] She uses another analogous example in ibid., 130-131.
[39] This is explicit in a number of places in Kamm's work, see for example ibid., 123, 130-131.

choose whether or not to pool their baseline chances for bad reasons. To see this, consider *2v3 Lifeboat*, except that now we have two additional outcome options, we can choose to save just A or just B. Now imagine that A is extremely racist and because B is a member of a race she hates passionately, she would rather choose *not* to pool her chances with B, even if that gives her a greater chance of being saved. Even if A hates B, this is not a good reason not to pool their chances to survive. This is a particular version of a general problem – if whether the baseline chances are pooled depends on actual claimants' consent, then both the individual and outcome group chances to receive the benefit can be adversely affected by morally irrelevant (or in this case, immoral) decisions on behalf of the claimants. One could avoid this objection by dropping the notion of consent.

This poses a new problem. On Kamm's original presentation, which claimants' baseline chances should be pooled was determined by which claimants consented to do so; however the new modification drops the need for claimant consent. The new modification is silent on how to pool chances when some outcome groups are proper subsets of larger outcome groups. Do we only focus on maximal outcome groups, where maximal outcome groups are outcome groups that are not proper subsets of any larger outcome groups? Or do we run the lottery between all outcome groups, even outcome groups that are proper subsets of maximal outcome groups?

For purposes of presentation, I will return to this problem in Section V.B after outlining Timmermann's individualist lottery as both the modified *proportional chances* and the individualist lottery are subject to another important objection.

III.B – Timmermann's Individualist Lottery

According to Timmermann's individualist lottery, every individual claimant should be given an equal chance to be selected by the lottery (again 1 divided by the number of equally worthy claimants) and the lottery is conducted without pooling any chances. However, once we have selected one claimant by lottery, we then benefit any other claimants we can while benefiting that selected claimant. That is, in *2v3 Lifeboat*, if we select A by the lottery procedure, we should also save B and if we select E we save C&D&E because we can save them together. According to Timmermann, this is because when we move to save A, we incur a new obligation to save B. Our previous obligation to give everyone appropriate and equal chances of being selected has been met and now that we can benefit B as well as A, we should do so. I will now outline an objection that applies to both Timmermann's and Kamm's weighted lotteries.

Both types of weighted lotteries discussed so far are unable to provide action prescription in cases where one or more claimants can be benefited in more than one outcome group. To see this, consider an equal conflict case where we have to choose to benefit only one of the following groups of claimants: A&B, A&C, D&E or D&F. This is a particular equal conflict case where because it is

possible to benefit some claimants (in this case, A and D) in more than one outcome group - that is, when outcome groups are not mutually exclusive. Call such cases *overlap cases*. *Unequal chance conflict cases* and *equal zero benefit cases* are particular kinds of *overlap cases*. The modified *proportional chances* procedure does not tell us how we should distribute chances of receiving benefit in cases like these. Similarly, if A was selected on Timmermann's lottery, it would not be clear whether we should benefit A&B or A&C[40]. To my knowledge, no philosopher has put forward a lottery procedure for distributing benefits in equal conflict cases that considers overlap cases, perhaps because the debate has focused on lifeboat style rescue cases where all outcome groups are mutually exclusive. As I will argue in the next section, there are many possible ways of extending the theory to cover these cases and the simplest options, such as iterating Timmermann's lottery, give implausible results in an important set of cases. When carefully examined, overlap cases present a great deal of complexity. This, in turn, necessitates significant complexity in the lottery procedures, as we will see in the next section.

The problem of overlap cases is not merely a technical concern, as many real life distribution cases have this feature. For example, there is often a group of patients all suffering from the same, equally bad disease, who all have a claim on treatment but different patients are receptive to different drug treatments, and the public health authority cannot afford all treatments[41]. Some patients are receptive to a number of different treatments, whereas others are not. That

[40] Thanks to Michael Otsuka for suggesting this point.
[41] Thanks to Dan Wikler for suggesting this problem from public health.

is, different treatments cure different patient outcome groups with overlapping members. One way this can arise is because of different blood type compatibilities amongst patients, with some patients being receptive to multiple blood types but others not. It is also becoming an increasingly relevant issue in personalized medicine as new genetic based treatments are developed that are only effective on people with certain gene structures and not others. Another way in which such a distributional problem may arise is in the production of a medical drugs - say drug A can cure a central case of disease X (e.g. squamous cell lung cancer) and two nearby variants of the disease, call them X* and X** (e.g. adenocarcinoma of the lung and adenosquamous cancer) and another drug B can cure disease X but a different variant of the disease, say X*** (e.g. small cell lung cancer), but capacity constraints prevent us from making both drug A and drug B. I discuss different ways of accounting for these more complex equal conflict cases in the next section.

Section IV – Introducing Composition Sensitive Lotteries

Once one considers equal conflict cases in which outcome groups are not necessarily mutually exclusive, vastly more lottery procedures become possible. When considering equal conflict cases with mutually exclusive outcome groups, as in standard lifeboat cases, the only plausible way of weighting chances of benefiting is by the size of the groups. Given all potential beneficiaries have equal claims to an equal benefit by stipulation, the size of each group is the only

plausible morally relevant difference. This is why although Timmermann and Kamm utilise different procedures, each claimant (and outcome group) has ultimately the same chance of benefiting on either procedure. Thus, when one only considers mutually exclusive outcome groups, there are practically only two types of lotteries, one that gives all outcome groups an equal chance of benefiting and one that gives each outcome group a chance of benefiting that is weighted according to its size.

However, when one considers overlap cases, in particular unequal chance conflict cases, only weighted lotteries can provide a lottery based account of all equal conflict cases and there are many more possible ways to weight a lottery. The focus on mutually exclusive lifeboat rescue cases has obscured the full complexity of the different kinds of possible lottery procedures. I now take on the taxonomical task of outlining and evaluating some of the new possible lotteries. I will call lotteries that are able to distribute chances of benefiting in these overlap cases *composition sensitive* lotteries, so called because they are sensitive to the composition of the outcome groups in a way that previous lotteries are not. Because of unequal chance conflict cases, all composition sensitive lotteries that account for all possible equal conflict cases are weighted lotteries. Therefore it turns out that the debate on which lottery it is fair to run in equal conflict cases ought not to be about whether one should adopt a weighted lottery or an unweighted lottery, as it has been in the past, but instead which composition sensitive weighted lottery is fair. I will begin exploring possible composition sensitive lotteries by first discussing different ways of

distributing chances in cases where claimants are in more than one outcome group. For presentational purposes and simplicity, I will only discuss cases where all outcome groups are maximal outcome groups at first. I will discuss the added complexity of non-maximal outcome groups later in Section V.B.

IV.A - Composition Sensitive Lottery Variants

In order to ensure that each equally deserving claimant is equally respected in the lottery procedure, comparative fairness requires each claimant to start with an equal baseline chance of 1/n that they contribute to the outcome groups of which they are a member. There are two possible ways in which they can contribute this baseline chance to the outcome groups of which they are a member - they can either divide their baseline chance equally amongst all such outcome groups or they can divide it unequally.

On the equal chances version, the chance contribution of any claimant *A* to an outcome group of which they are a member is $(1/n)/O_A$, where *n* is the total number of claimants and O_A is the number of outcome groups of which claimant *A* is a member. Call this the *equal composition sensitive lottery*. This lottery captures the idea that from any individual claimant's perspective, each of the outcome groups they are a member of are equally important as they are benefited in each of them. I will provide an example of this chance distribution in the upcoming section. The chance of benefiting a particular outcome group, M, is given by the following formula:

$$P_{\text{equalcomp}}(M) = \frac{1}{n} \sum_{x \in M} \frac{1}{O(x)}$$

Other composition sensitive lottery variants do not distribute individual claimant's baseline chances equally amongst all outcome groups, but instead weight those contributions. I can see two independently motivated unequal ways of distributing a particular claimant's baseline chance between the different outcome groups of which they are a member, one I will call the *inclusive (composition sensitive) lottery* and the other, the *exclusive (composition sensitive) lottery*.

On the *inclusive lottery*, a claimant A would contribute a w_Y/t proportion of their baseline chance to group Y, where w is the number of claimants in group Y of which A is a member (i.e. the size of group Y) and t is the sum of *all of the sizes* of the outcome groups of which A is a member. So A would contribute $(w_Y/n)/t$ to each outcome group of which it is a member, with w changing depending on the size of the particular outcome group under consideration. This would weight each claimants' individual chance contributions to the different outcome groups of which they are a member on the basis of the sizes of the outcome groups. The inclusive lottery captures the idea that it is the absolute size of each outcome group that matters and that absolutely larger groups should receive a proportionally larger amount of the baseline chance. It should be noted that when the claimant is only in one outcome group, w=t and therefore

they contribute their entire baseline chance to the sole group of which they are a member.

This is best illustrated by an example. Consider a case in which we have to distribute a benefit to A&B&C or C&D. In this case, each claimant begins with a 1/4 (25%) baseline chance because there are four claimants in total. A and B each contribute all of their baseline chance (25%) to A&B&C. C contributes 3/5 of their baseline chance (1/4 x 3/5 = 3/20 = 15%) to A&B&C and 2/5 of their baseline chance to C&D (1/4 x 2/5 = 2/10 = 10%). D contributes their entire baseline chance (25%) to C&D. A&B&C therefore receives a 65% chance of benefiting (25% + 25% + 15%) and C&D receives a 35% chance of benefiting (10% + 25%).

If one wanted to take into consideration not simply the total size of the outcome groups, but instead their composition, one could adopt what I call the *exclusive lottery*. To explain this lottery, I will have to introduce a new term. When considering a set of outcome groups of which a claimant A is a member, an *exclusive* claimant is any claimant that is in at least one but less than all of the outcome groups of which A is a member. That is, a non-exclusive claimant is one that is in *all* or *none* of the outcome groups of which A is a member. Therefore whether a claimant is exclusive or non-exclusive depends on which claimant we are considering. A claimant can be exclusive relative to one claimant's outcome groups but not another's. For example, consider the case when you have to benefit A&B&C&D or A&B&C&E or A&F or B&F. From C's perspective, A, B and C are not exclusive as they can be benefited in all of the

groups that C can be benefited in. However, from F's perspective, A & B are exclusive as they are in some but not in all of the outcome groups that F can be benefited in.

On the *exclusive lottery*, a claimant A contributes a u_Y/e_A proportion of their baseline chance to group Y of which A is a member, where u is the number of *exclusive* claimants in group Y and e is the total sum of the number of *exclusive* claimants in *all of* the groups of which A is a member. This exclusive lottery captures the idea that what matters in the distribution of baseline chances are the compositional *differences* between the outcome groups of which each claimant is a member, not merely their size. By excluding non-exclusive claimants, people who can be benefited in all groups alongside a particular claimant are ignored in the calculation of how that particular claimant's baseline chance contribution is divided amongst the outcome groups of which they are a member. That is, people who can be benefited in all possible options alongside the claimant who is under consideration are ignored as they are benefited in all such possible options. This would weight each claimants' individual chance contributions to the different outcome groups of which they are members depending on how those groups are composed (specifically, how many exclusive claimants each group contains). When the claimant is only in one outcome group, u=e=0, as all other claimants in that outcome group are not exclusive by definition. However as dividing by zero is incalculable, we need to determine how to distribute chances in these cases in a theoretically motivated way. As the claimant is only in one outcome group, for the purposes of the

lottery they should contribute their entire baseline chance to the outcome group of which they are a member. The exclusive lottery is appealing as it is sensitive not only to relative outcome group size but more importantly, the composition of those groups.

To illustrate, consider a case in which we have to distribute a benefit to A&B&C or A&B&D&E. Each claimant begins with a 1/5 baseline chance because there are five claimants in total. When calculating A's chance contributions to the outcome groups, group A&B&C contains one exclusive claimant (C), whereas group A&B&D&E contains two exclusive claimants (D&E) and total sum of the number of exclusive claimants in *each of* the groups of which A is a member is three (A&B are not exclusive). Thus, A contributes 1/3 of their baseline chance to A&B&C (1/3x1/5=1/15) and 2/3 of their baseline chance to A&B&D&E (2/3x1/5=2/15). Similarly, as B's contribution also ignores non-exclusive claimants A&B, then B contributes 1/3 of the baseline chance to A&B&C (1/3x1/5=1/15) and 2/3 of their baseline chance to A&B&D&E (2/3x1/5=2/15). C contributes their entire baseline chance (1/5) to A&B&C. D&E contribute their entire baseline chances (1/5 each) to A&B&D&E. A&B&C therefore receives a 1/3 (1/15+1/15+1/5) chance of benefiting. A&B&D&E receives a 2/3 chance of benefiting (2/15+2/15+1/5+1/5).

IV.B - A Composition Sensitive Variant of Timmermann's Individualist Lottery

It is possible to modify Timmermann's individualist lottery to account for cases in which some claimants can be benefited in more than one outcome group. On such a modified view, after the first lottery is run, Timmermann could run multiple iterated lotteries between outcome groups that can be benefited alongside the previous lottery winner(s). These iterated lotteries would determine who to benefit in complex equal conflict cases in which claimants belong to multiple outcome groups. For example, imagine that we could benefit A&B&D, A&B&E, A&C&F, A&C&G or H and the first lottery winner is A. After A has been selected, one would have to run the individualist lottery between B,C,D,E,F and G (those who can benefited with A). If the second lottery winner is C, then a third lottery would then have to be conducted to see who to benefit out of F and G (those who can benefited with A & C).

Now that I have completed my taxonomy of all of the composition sensitive lotteries for which I can provide independent motivations, which lottery should we adopt as the fair procedure? None of the independent motivations for each of the different lotteries gives us decisive reason to adopt one over the other, so I will now utilise a case based evaluation.

Section V - Evaluating Composition Sensitive Lotteries

V.A - Convergent Cases

One way of distinguishing the plausibility of the different composition sensitive lotteries is by comparing the chances of benefiting they distribute with our moral judgments of the fairness of those distributions. After all, these lotteries are supposed to be the most fair way of distributing the benefits in equal conflict cases. However, all weighted composition sensitive lotteries so far discussed prescribe precisely the same chances of benefiting in cases where all outcome groups are mutually exclusive and prescribe very similar chances of benefiting in cases where outcome groups only share a small proportion of their members. Therefore, we should look to more extreme limiting cases, where the distribution of chances to benefit will diverge. One important class of such cases are what I will call *convergent cases*. *Convergent cases* are equal conflict cases in which different outcome groups have a large proportion of the same members. For a particularly difficult convergent case, consider 1000 people, numbered sequentially. One outcome group, O_1 contains claimants numbered from 1 to 500 (500 claimants), group O_2 contains claimants numbered 501 to 1000 (500 claimants) and group O_3 contains claimants numbered 2 to 999 (998 claimants)[42].

[42] Both Richard Chappell and Andrew Schroeder have independently suggested that one could adopt maximin of the chances of benefiting as a method of distributing chances that is informative in *unequal chance conflict cases*. However, such a distribution of chances has many problems. Firstly, like unweighted lottery views, it violates *Sensitivity* and therefore does not take into account an intuitively morally relevant feature that I motivated in Section II. Secondly, it provides implausible and unfair chance distributions in many cases. For example, in the above

The *equal composition sensitive lottery* will allocate the following chances of benefiting: 25.05% (1/1000 + 499x1/1000/2) to O_1, 25.05% (499x1/1000/2 + 1/1000) to O_2 and 49.9% (998x1/1000/2) to O_3. This gives the startlingly counter intuitive result that it is fair to give a greater than 50% chance to save either O_1 or O_2, both of which contain approximately half as many members as the largest group. According to the *equal composition sensitive lottery* in this case, it is more likely that we will save 500 people than 998, which is highly implausible. This happens because the *equal composition sensitive lottery* distributes the baseline chances of each claimant *equally* amongst all outcome groups of which they are a member and does not take into account either the size of the outcome groups nor how similar (in terms of membership) they are when doing so. Note that it is no objection to these procedures to claim that they provide excessive precision in fractions of percentages when prescribing chances of benefiting.

---

convergent case, maximin will give 50% chance of benefiting to 1-500, and 50% to 501-1000 and 0% chance to the biggest group of most members. It would give even more implausible chance distributions in other cases, where again it would give no chance to the biggest outcome groups to be benefited. For example, imagine you can benefit one of 1&2, 3&4, 5&6.... 99&100, or the set containing all the even numbered individuals. Maximin would give a 2% chance of benefiting each pair, 100% of benefiting some pair, and 0% chance of benefiting the 50 even numbered invididuals. Note these prescriptions are the same as are provided by unweighted lotteries and thus provide objections to unweighted lottery theorists. Thirdly, using maximin to determine chances of benefiting will also conflict with a particularly plausible version of maximin at the level of actual benefits distributed. Maximin is concerned about making the worst off as well off as possible, so it orders possible distributions from best to worst, regarding maximin, on the basis of how well off the worst off are. In these scarce benefit distribution cases, the worst off are always at the same level, as they do not receive the benefit they have a claim on. In such cases, it is plausible to order distributions according to a version of maximin that firstly (lexically prefers) maximises how well off the worst off are, then minimizes the size of the worst off group (i.e. tries to improve some of the worst off so they are no longer the worst off). In the O_1 O_2 O_3 case above (and many other cases), when maximin at the chance distribution level is adopted, it will result in more worse off people when considering the actual benefits distributed. In the first convergent case, it will have the result that there are 500 people in the worst off position, namely those that haven't received a benefit. This may be a very bad result if the benefits are great, as they are in life rescue cases. Much better off, by maximin of actual benefits, understood in this way, to give the benefit to 2-999. For these three reasons, using maximin to determine the chances of benefiting is implausible.

All procedures in the literature give precise chances of benefiting which are sometimes complex in virtue of composition of outcome groups. Even the simplest unweighted lottery view prescribes precise chances of benefiting in some complex cases. In a case where you can just benefit any one of forty six claimants, an unweighted lottery prescribes that all claimants receives a 2.17391...% of benefiting.

The *inclusive lottery* will distribute chances of benefiting as follows. Individuals 2-500 would each contribute (500/1000)/1498 to O_1 and (998/1000)/1498 to O_3. Similarly, individuals 501-999 would each contribute (500/1000)/1498 to O_2 and (998/1000)/1498 to O_3. Individuals 1 and 1000 would contribute their entire 1/1000 baseline chance each to O_1 and O_2 respectively. The resulting chances of benefiting would be approximately 16.8% each for groups O_1 and O_2, with the largest chance of approximately 66.4% of benefiting the largest group of O_3. This is a more intuitively palatable result than the *equal composition sensitive lottery,* but it still gives an implausible and intuitively unfair high chance of benefiting of over a third (approximately 33.6%) that we will save 500 people rather than 998. This is because while the *inclusive lottery* takes into account the size of the outcome groups when distributing baseline chances between outcome groups, it does not take into account how similar (in terms of membership) they are. Given the similarity in the memberships of the overlapping outcome groups, one would expect the exclusive lottery to fare better. Let's see.

The *exclusive lottery* will distribute chances of benefiting as follows. Individuals 2-500 would each contribute (1/1000)/500 to O_1 and (499/1000)/500 to O_3, as there is one *exclusive* claimant in O_1 and 499 *exclusive* claimants in O_3. Similarly, individuals 501-999 would each contribute (1/1000)/500 to O_2 and (449/1000)/500 to O_3. Individuals 1 and 1000 would contribute their entire 1/1000 baseline chance each to O_1 and O_2 respectively. The resulting chances of benefiting would be approximately 0.2% each for groups O_1 and O_2 ((1/1000)/500x499 + 1/1000), with the largest chance of approximately 99.6% of benefiting the largest group of O_3 ((499/1000)/500x499+ (499/1000)/500x499). This gives by far the highest chance of benefiting to the group that contains the vast majority of all claimants, which is intuitively fair.

Working out the chances of benefiting each group using my modification of Timmermann's lottery is more complex. I will first explain O_1's chance of being benefited. The same remarks apply, *mutatis mutandis*, to O_2. As should become clear, the chance that O_1 is benefited is the same as the chance that O_2 is benefited. The chance that O_1 will be benefited on the first lottery is the chance that 1 will be selected out of the 1000 possibilities (1/1000). Then we have to work out the chance that O_1 will be benefited in subsequent lotteries. This is the chance that 1 will be selected after one of 2-500 is selected. The sum of the chances of this happening is as follows:

(499/1000)*(1/998)+(499/1000)*(498/998)*(1/997)+(499/1000)*(498/998)*(497/997)*(1/996)+...

The first term corresponds to the chance that someone from 2-500 will be picked on the first draw, followed by number 1 on the second draw. For O_1 to be selected on the second draw, a number between 2-500 has to be selected on the first lottery (a 499/1000 chance) and then 1 has to be selected by the second draw (a 1/998 chance, 998 because the claimant in O_1 that was selected by the first lottery is removed from all subsequent lotteries and claimant 1000 can no longer be benefited). The second term corresponds to the chance that someone from 2-500 will be selected on the first two draws followed by number 1 on the third draw, and so on. Summing over all possible ways the lotteries could go up to 500 iterations, we obtain the following expression of the chance O_1 will be selected (the factorial fraction in the middle is the chance that 1 has to be selected on the unlikely chance of a 500th lottery):

$$P(1-500) = P(501-1000) = \frac{1}{1000} + \frac{1}{1000}\sum_{n=1}^{499}\frac{499!}{(499-n)!}\frac{(998-n)!}{998!} \approx 0.2\%$$

The chance of O_3 being benefited is the sum of the chances that someone from 500-999 is selected after someone from 2-500 and the chances that someone from 2-500 is selected after someone from 500-999 in any of the iterated lotteries. This is expressed below:

$$P(2-999) = \left[\frac{998}{1000}\sum_{n=1}^{498}\frac{499^n}{999!}(999-n)!\right] - \frac{2}{1000} \approx 99.6\%$$

The 2/1000 at the end of that equation is the chance that either 1 or 1000 is selected on the first lottery, thereby preventing O_3 from being benefited.

Like the *exclusive lottery*, the composition sensitive version of Timmermann's lottery plausibly gives a high chance (99.6%) of O_3 being benefited. How ought we to pick between these two lotteries that give plausible results so far? I turn to answer this question by now adding the additional complexity of non-maximal outcome groups.

V.B - Non Maximal Outcome Groups

What is the fair distribution of chances when some outcome groups are proper subsets of other outcome groups? My iterated version of Timmermann's lottery readily has a solution for this, as it can continue running lotteries until all lottery winners are members of only one maximal outcome group. However the question is more difficult for the *exclusive lottery*. Returning to the question from Section III.A, should baseline chances be distributed between all outcome groups or only maximal outcome groups?

In answering this question, it is useful here to distinguish two important components of *all things considered* fairness in these cases. There is *comparative* fairness, which requires all equally deserving claimants to be treated equally in the outcome group selection procedure and *absolute* fairness, which is best satisfied by actually giving all claimants the benefit they have a

claim on[43]. As we saw in the *Chancy Doctor* case, the fairest situation is for everyone to *actually benefit* while getting a 100% chance of benefiting, and the most fair procedure in cases of scarcity will best approximate this result. As a matter of *comparative* fairness between groups, it does not matter that some outcome groups are proper subsets of other groups. Every group of equally deserving claimants is due their appropriate chance of benefiting. However, giving the largest outcome group 100% chance of benefiting will satisfy the most *ex post* claims, but at the expense of *comparative* fairness as equally deserving claimants will not only be treated unequally, they will be given inappropriate chances of benefiting. How do we address these two competing concerns? My iterated variation of Timmermann's lottery provides guidance here.

As it is irrelevant to *comparative* fairness whether or not outcome groups are maximal or not, the exclusive lottery ought to distribute claimants' baseline chances between all outcome groups, not merely maximal outcome groups. However, in order to be sensitive to *absolute ex post* fairness concerns, when a non-maximal outcome group is selected by the lottery, we can benefit the larger maximal outcome group of which it is a proper subset. This is, after all, the same motivation behind the move to save whoever can be saved alongside the winners of Timmermann's lottery. There is just one more complication we have to consider for the exclusive lottery. What is the fair thing to do when there is more than one maximal outcome group of which the selected non maximal

[43] As we saw in the preceding sections, equally fair treatment is *not* achieved by giving everyone an equal chance of benefiting but an equal baseline chance, as that gives them all an equal contribution to the outcome group selection procedure.

group is a proper subset? Again, my modified version of Timmermann's lottery gives us guidance – we can just iterate the exclusive lottery to decide between the different maximal outcome groups. That is when one non-maximal outcome group that is a proper subset of more than one maximal outcome group wins an exclusive lottery, a further exclusive lottery is run between all of the different outcome groups that include the proper subset that won previously. This continues until a maximal outcome group has been selected.

Consider the following illustrative example. Imagine you can benefit A&E, B, A&B&C, or A&B&D&E and non maximal outcome group B is selected first. The maximal outcome groups of which B is a proper subset are A&B&C and A&B&D&E. We now treat these two maximal outcome groups as the only options in a new exclusive lottery. There are five claimants, so each has a 1/5 baseline chance. From A's perspective, C, D & E are exclusive claimants, one in the former group, two in the latter group. A therefore contributes 1/3rd of their baseline chance to A&B&C (1/3 x 1/5 = 1/15) and 2/3rds of their baseline chance to A&B&D&E (2/3 x 1/5 = 2/15). Similarly, from B's perspective, B contributes 1/3rd of their baseline chance to A&B&C (1/3 x 1/5 = 1/15) and 2/3rds of their baseline chance to A&B&D&E (2/3 x 1/5 = 2/15). C can only be benefited in the former group, so C contributes their entire baseline chance to the former group (1/5). D & E can only be benefited in the latter group, they each contribute their entire baseline chance (1/5) to the latter group. Thus A&B&C gets a 1/3 chance of benefiting (1/15+1/15+1/5) and A&B&D&E get a 2/3 chance of benefiting (2/15+2/15+1/5+1/5).

In order to decide which of the two remaining procedures is most fair, I shall now consider a complex *convergent* case with non-maximal outcome groups. As before, consider 1,000 claimants, numbered consecutively. The outcome groups we can benefit are 1-500, 501-1000 and every possible pair. According to Timmermann's lottery, the chance the first group of 500 will be selected is the chance of a claimant from 1-500 being selected on the first lottery (1/2) multiplied by the chance of selecting one of the 499 remaining claimants of that group (499/999), which is approximately 24.975%. This is the chance of selecting a non-maximal pair in the first group, which results in benefiting the maximal outcome group of which they are a member. There is the same chance of benefiting the second group of 500. The chance of some maximal outcome group of two being selected is the same as the chance that neither group of 500 is selected – just over 50%. Thus, according to Timmermann's lottery, the most likely option is that an outcome group containing just two members will be ultimately benefited, which gives an excessively high, unfair chance to these smaller groups of two. Let us consider how the exclusive lottery does.

According to the exclusive lottery, from any claimants' perspective, there are 499 exclusive claimants in 1 outcome group, and 1 exclusive claimant in each of 999 outcome groups to distribute their 1/1000 baseline chance between. 499 of those 999 pairs will be proper subsets of one of the maximal outcome groups of 500, so if they are selected, then one of the largest outcome groups will be selected. The chance of the outcome group of 500 being benefited is 500 x

1/1000 x (499+499)/1498 is approximately 33.3%. Thus on the exclusive lottery, the chance of one of the two groups of 500 of benefiting is a significant majority of approximately 66.6%. There is a still substantial chance to benefit one of the smallest groups of two. This occurs because the procedure respects each and every claimants' equal claim by giving each of them an equal impact on the chancy outcome group selection procedure, while giving all outcome groups a proportional chance of benefiting. From this convergent case with non-maximal outcome groups, I contend that the exclusive lottery is the most fair procedure in equal conflict cases. However, while such a procedure is fair, its prescriptions are not necessarily right – it would be deeply wrong to save 2 when one could save 500 instead. I now deal with this issue with a discussion of all things considered fairness and the importance of fairness to what is morally right.

Section VI - The Relation of the Fair to the Right

Before discussing the relation of the fair to the morally right, it is important to be clear how the *exclusive lottery* is the *all things considered* fair procedure. It is *all things considered* fair because it best satisfies the sometimes conflicting requirements of both *comparative* and *absolute* fairness. Comparative fairness requires all equally deserving claimants to be treated equally. Equality of what kind of treatment? In Section II I argued that comparative fairness does not, as Broome contends, require all claimants to receive equal chances of benefiting in equal conflict cases. Instead, I argued that comparative fairness requires all

claimants to have an equal impact on the *ex ante* outcome group selection procedure in a way that respects the *Sensitivity* principle. This requirement is met by Kamm's proportional chances, Timmermann's individualist lottery, and the *composition sensitive* variants of those views proposed in Section IV and V. However both Kamm's & Timmermann's original procedures do not account for *overlap* cases and should therefore be rejected.

Absolute fairness is only achieved when people are actually given what they are due, which in these cases is when claimants actually receive the benefit. *Absolute* fairness requires as many people as possible to actually receive the benefit *ex post*. When benefits are scarce and indivisible as they are (by stipulation) in equal conflict cases, then the requirements of *comparative* fairness and *absolute* fairness can conflict[44]. However there is a reasonable compromise between these conflicting requirements. Of the *comparatively* fair lotteries, the *exclusive lottery* gives an intuitively plausible, relatively high proportion of the chances of benefiting to the largest maximal outcome groups. The *exclusive lottery* is therefore the *comparatively* fair lottery that best satisfies *absolute* fairness. This lottery was not manufactured *ad hoc* to meet this requirement of *absolute* fairness. The *exclusive* lottery has this feature because it gives all equally deserving claimants an equal baseline chance that they distribute between outcome groups of which they are a member according to

[44] They do not conflict, for example, in the aforementioned *33v33v33 Train Cars*. All options will satisfy the same number of claims *ex post*, so all comparatively fair options will best satisfy the requirements of absolute fairness. The most common case in which they conflict are when there are mutually exclusive outcome groups that are of unequal size. For example, they *do* conflict if the numbers are changed so the situation becomes *33v33v32 Train Cars*.

the differences in the outcome groups' size and composition. Furthermore, we have independent reason from the discussion in Sections II & III that it is the *differences* in group size and composition that are relevant to the fair chances of benefiting. Thus the *exclusive lottery* best satisfies the sometimes conflicting requirements of both *comparative* and *absolute* fairness. While the *exclusive* lottery is the all things considered fair procedure, is it all things considered morally right? To this question I now turn.

It is beyond the scope of this paper to provide a complete account of how we should distribute benefits in all equal conflict cases, but I will offer a sketch of how procedural fairness affects what we ought to do in such cases. Both consequentialists, such as standard act utilitarians, and nonconsequentialists, such as Kamm[45] and Scanlon[46], claim that we always ought to benefit the greatest number in cases such as this. In order to cut across this familiar divide, I will call theories that claim we always ought to benefit as many people as possible in equal conflict cases *majoritarian theories*. If we ever ought to act according to my account of fairness, my account of the fair *ex ante* chances of benefiting conflicts with majoritarian theories. This is because majoritarian theories claim we ought to give 100% chance of benefiting to the largest outcome group in equal conflict cases[47]. Should we ever act according to my account of fairness?

[45] ibid., 116-119.
[46] Scanlon, *What we owe to each other,* 229-234.
[47] That is, assuming that there is only one maximal outcome group which is largest. In a case where there are a number of maximal outcome groups which all tie for largest size, someone who wanted to always benefit the greater number would have to distribute all 100% chance of

I contend that we should do so in some cases - to see this, consider *10,000v10,001 Cruise Ship*[48]. This case is identical to *Lifeboat* except that there are much larger numbers of lives at stake. In this equal conflict case, it does not seem straightforward that one should just save the greater number and give no chance for the 10,000 to be saved. Not only do majoritarian theories suggest that we ought to save the larger number, condemning the other 10,000 to death with no chance of survival, they objectionably appear to regard such an action as *obviously* the right answer[49]. To give each Cruise Ship approximately 1/2 chance of surviving according to the exclusive maximal lottery is fairer than giving 10,000 no chance of survival and is intuitively also the right action to take.

However, we ought not always run a lottery to distribute chances. While I maintain that the exclusive maximal lottery is the fair method of distributing chances of benefiting in equal conflict cases, it is not the only morally relevant consideration and thus what is fair is not co-extensive with what is right[50]. The

---

benefiting between these equally large groups. Kamm has suggested that in such cases we ought to distribute between all groups equally. Whereas Scanlon has suggested in such cases it would not matter which of the largest maximal outcome groups we ought to select as long as we did not do so based on some arbitrary and morally irrelevant feature. Still, my account of fairness conflicts with these cases as well, as I give some chances of benefiting to smaller maximal outcome groups.

[48] This case is similar to a case presented by Kamm which she attributes to Peter Unger. See Kamm, *Morality, Mortality: Volume 1,* 103, 121n4.

[49] This echoes a criticism made of utilitarianism by Bernard Williams throughout his contribution to John Smart and Bernard Williams, *Utilitarianism; for and Against,* Cambridge University Press, 1973)

[50] My weighted lottery view is therefore different from views such as Dan Brock's that advocate weighted lotteries as the morally right compromise between fair equal chances of benefiting and the value of the actual benefits distributed. According to such views, running a weighted lottery is fairer than running no lottery at all, but not as fair as running a lottery that gives all an equal chance of benefiting. While Brock has defended this view in print, in personal conversation with me he suggested that he may revise his views in response to my arguments.

exclusive lottery is fair because it gives all *equally* deserving claimants an *equal* baseline chance that they distribute appropriately amongst all outcome groups of which they are a member. However, fairness has to be balanced against the potential good consequences achieved by different benefit distributions in order to establish what the right benefit distribution is. After all, remember from Section I that procedural fairness is by definition the appropriate way to satisfy *claims* and claims are only a particular subset of all reasons (namely reasons that are owed to a particular agent or group)[51]. Fairness is not the sole source of reasons. The value of the benefits distributed also provides reasons.

It is intuitively wrong to run a benefit distributing lottery in *1v10,000 Cruise Ship* that gives any chance for the one to be saved[52]. Even if it allocates a very small chance of benefiting a sole individual, it would be morally wrong to save that one person even if they were selected by that fair lottery. To represent the balance between fairness and the good consequences achieved by distributing benefits, I advocate implementing what I call the quantitative moral relevance principle. On my view, the most fair procedure in equal conflict cases is to run my exclusive composition sensitive lottery. The procedure with the best

---

For his views in print see Brock, Dan 'Ethical Issues in Recipient Selection for Organ Transplantation' in D. Mathieu, *Organ substitution technology: ethical, legal, and public policy issues,* Westview Pr, 1988)

[51] To see an example of this, compare a case where you have promised two of your friends to give them something of no value to a case where you have promised those same friends to look after their young children for a week while they are away and no one else can look after the children. In both cases, the promises generate claims and there are reasons of fairness to give your friends what they are due. However in the latter but not the former case, there are reasons *in addition* to the claims in virtue of the importance of looking after those children.

[52] Cases like these are used as objections to using lottery distribution procedures. I agree that it is not morally right to run a lottery in such cases, but maintain that it is still fair to do so. For versions of this objection, see Iwao Hirose, "Review Article: Aggregation and Non-Utilitarian Moral Theories," *Journal of Moral Philosophy* 4, no. 2 (2007): 273-284.; Hirose, *Weighted Lotteries in Life and Death Cases,* 45-56.

consequences (excluding considerations of procedural fairness) is usually[53] a majoritarian one. My *quantitative moral relevance principle* strikes a balance between these two when potential good consequences can justify some unfairness. My *quantitative moral relevance principle* states that:

> when good consequences justify some unfairness, then outcome groups that will produce significantly worse consequences if benefited should be treated as irrelevant in the exclusive lottery and thus not allocated any chance to benefit.

An outcome group produces significantly worse consequences for the purposes of quantitative moral relevance when benefiting that outcome group produces consequences that are substantially bad enough when compared to the option with the best consequences that they outweigh the value of fairness gained by including that outcome group in the composition sensitive lottery. Applied to life rescue cases, because potential beneficiaries' claims are defeasible and peoples' lives are extremely morally valuable, the value of peoples' lives are, in many cases, defeaters. This principle recognizes both fairness's importance and defeasibility.

[53] There are some rare cases in which the best consequences can only be brought about by *not* directly benefiting the largest outcome group, for example when benefiting a smaller group enables that smaller group to benefit others (e.g. the situation is *1v2v3 Lifeboat*, but the group of one is a lifesaver who can save three other claimants' lives). In such cases, we ought to treat the lifesaver and the people she can benefit as an outcome group even if we do not provide the benefit directly to the claimants. Thus the value of the consequences brought about is not always the benefits being directly distributed.

Sometimes we ought to run a lottery that is not as fair as it could be in order to recognize the moral importance of other reasons. To see this, consider an additional cruise ship is added to the case, as in *1v10,000v10,001 Cruise Ship.* Applying my *quantitative moral relevance principle* you should run a lottery that does not consider the isolated individual and gives a 10,000/20,001 chance to save the 10,000 and 10,001/20,001 chance to save the 10,001. This procedure gives reasonably fair chances to the vast majority of claimants, while respecting the importance of lives saved. Therefore sometimes we ought to run a lottery that excludes some outcome groups, even if it is not the most fair procedure. This ensures some level of fairness. In other cases we ought just to give the largest number 100% chance of benefiting.

Furthermore, because our reasons from fairness can be weighed against the value of the benefits distributed, I argue that is also sometimes right to engage in my weighted lottery. This is so even if *some* reduction of the benefit conferred to claimants is necessary for the weighted lottery. For example, imagine

> *99v100 Swimmers,* where all claimants are currently treading water to prolong their lives before drowning. Our three deliberative options are to immediately save the larger group of a hundred, immediately save the smaller group of ninety nine or engage in a composition sensitive lottery to select which group to save that takes time so that after doing so, the group we choose to save has had to endure the pain of having to tread water for

two minutes before being saved (the group that drowns would have to endure that pain no matter what option is selected as they attempt to prolong their life).

Choosing immediately to save the greater number instead of engaging in my composition sensitive lottery procedure avoids having the larger group of a hundred endure this pain. My intuition is that despite this increased pain, a weighted lottery is still the right and fair thing to do. As such, this intuition supports my earlier claim that sometimes the right thing to do is not that which produces the best outcomes and that we should engage in a composition sensitive lottery even when doing so leads to somewhat worse outcomes than choosing to benefit the largest group of claimants. However, as before, I should stress that I am not claiming that we should always engage in composition sensitive lotteries to select whom to benefit. Composition sensitive lotteries should not be conducted at all costs. It is not right to engage in a weighted lottery if doing so reduces the chance of saving the individuals we select from 100% to 5%. That is, if our three deliberative options are to immediately save the group of a hundred, immediately save the group of ninety nine, or engage in a composition sensitive lottery that takes time so that after doing so, we only have a 5% chance of saving the people in the group we select by lottery, we should save the hundred instead of conducting a composition sensitive lottery. This is because considerations of fairness are outweighed by considerations of outcomes. Fairness can motivate some reduction in goodness of outcome, but is not lexically prior to goodness.

I have now outlined my view about what distribution is fair in equal chance conflict cases and the relation between fairness and moral rightness. In the next section I will defend my lottery based view from three prominent objections. Firstly, Hirose argues that weighted lotteries are under motivated. Secondly, Harsanyi argues that prior 'natural' lotteries obviate the requirement to run subsequent lotteries. Thirdly, there will be very few actual cases in which the benefits distributed are exactly qualitatively equal and if lotteries are only to be used in breaking ties, the value of lotteries as a deliberative procedure (and relatedly, the value of fairness) would be extremely limited.

Section VII - Defending Lottery Procedures

VII.A - Against Pooling

Hirose objects to any weighted lottery procedure that involves pooling by arguing that individuals should not be allowed to combine their chances to receive benefits.[54] As such, his criticism applies alike to both Kamm's *proportional chances* and my composition sensitive lottery. To apply this to 2*v*3 *Lifeboat*, Hirose denies that Alice and Betty should be allowed to pool their chances so they have 2/5 chance of being saved. Hirose argues that because the only motivation for the lottery procedure is the equal division of chances to benefit each individual, if Betty is selected by the non-pooled lottery procedure

[54] ibid.

we should save Betty and Betty alone. Saving Alice as well as Betty does not comply with the precisely equal division of chances.

There are two problems with Hirose's objection. Firstly, Hirose, like many critics of the weighted lottery procedure, fails to appreciate all of the distinct motivations of weighting chances of benefiting. My composition sensitive lottery, unlike Broome's unweighted lottery, is not motivated by the equal division of chances to receive the benefit. My composition sensitive lottery accounts for intuitions supporting an outcome group selection procedure that does not always save the greater number in cases such as *10,000v10,001 Cruise Ship*. Furthermore, the fact that the weighted lottery *does* allow for pooling of chances allows it to respect *Sensitivity*. That is, it is sensitive in a way that unweighted lotteries are not to the morally relevant difference between *1v2 Lifeboat* and *1v100 Lifeboat*. Furthermore, unlike Taurek's unweighted lottery, my procedure is also sensitive in the appropriate way to addition of new outcome groups, for example when *33v33 Train Cars* is compared with *33v32v1 Train Cars*. This is because the outcome group that does not change in claimant composition (the second group of 33) when no new claimants are added to the situation (a total of 66 claimants in both cases) has the same chance of being saved in both cases and all groups are allocated a chance of survival directly proportional[55] to the number of equally worthy claimants they contain (thereby respecting group fairness). Allowing different individuals' chances of receiving benefits to be pooled gives more intuitively fair results in these cases because it

[55] The outcome group chance to benefit is directly proportional to number of equally worthy claimants in equal conflict cases when no claimants are in more than one outcome group.

respects both individual fairness (by giving each claimant an equal contribution to the outcome group selection procedure) and group fairness (by respecting the claimant composition of maximal outcome groups). This is related to my second response to Hirose.

Secondly, Hirose fails to acknowledge the moral importance of the distinction between individual baseline chances of receiving the benefit and chances of selecting a particular benefit distribution to act upon. Each worthy recipient is given an equal baseline chance of receiving the benefit (1/number of worthy recipients) thereby respecting our intuitions regarding individual fairness. When some worthy claimants' chances are pooled, it is not that some individuals are arbitrarily given a higher chance of receiving the benefit, it is because the structure of possible outcomes determines the different ways in which the benefit can be distributed. In this way, my composition sensitive lottery procedure is sensitive to both group fairness and individual fairness. In *2v3 Lifeboat*, one can either save Alice and Betty or Celine, Diane and Edna. My composition sensitive lottery procedure does not claim that any group of individuals can choose to pool – that would be objectionable. Chances of receiving benefits should be pooled only amongst worthy recipients who are all capable of receiving the benefit. That is, what determines whether or not individual chances can be pooled are the ways in which the benefit can be distributed, not an unequal distribution of chances to individuals. Thus pooling equal baseline chances respects individual equality while being sensitive to the different maximal outcome groups available.

VII.B - The 'Grand Lottery of Life' Objection

A number of philosophers including Harsanyi object to people conducting lotteries to decide who to benefit because they claim that there has already been a prior 'natural' lottery which obviates the need to run a new lottery[56]. This is often combined with the view that we ought to benefit the largest outcome group in equal conflict cases. By prior 'natural' lottery, they are referring to the fact that things could have been different in a relevant way, namely that potential beneficiaries had a chance to be in a different outcome group to the one they are currently in. For example, things could have been different than 33v33 Train Cars. Some of the people in the first train car could have been in the second train car, and vice versa. Because the potential beneficiaries had a chance to be in the outcome group that will be benefited (often, on their view, the largest outcome group), we have no reason to run another lottery when deciding who to benefit. Call lotteries run by people to distribute chances of benefiting 'non-natural' lotteries. Harsanyi rhetorically asks "Why should we assign higher moral dignity to a lottery organized by government bureaucrats than we assign to the 'great lottery of life'... ? Why should a bureaucratic lottery be regarded as being a 'fairer' allocative mechanism than the great biological lottery produced by nature?"[57]. Harsanyi's objection applies to both unweighted and weighted lotteries. Even though I am

[56] Julian Savulescu has raised this objection with me in conversation. Scanlon appears to express a view of this kind in Scanlon, *What we owe to each other,* 420. Harsanyi explicitly makes this objection against fair lotteries in John Harsanyi, "Nonlinear Social Welfare Functions: Do Welfare Economists Have a Special Exemption from Bayesian Rationality?" *Theory and Decision* 6, no. 3 (1975): 311-332.
[57] ibid.: 317.

an advocate of weighted lotteries, my response to Harsanyi can be utilized by both weighted and unweighted lottery theorists in defence of their views. In order to apply to both kinds of lottery views, I will focus on a case where all outcome groups are equally sized and mutually exclusive. In those cases, all lottery views considered in this paper agree that fairness requires us to give all outcome groups equal chances of benefiting.

While I agree that in many equal conflict cases, potential beneficiaries could have been in a different outcome groups than they actually are, I will argue that this alone does not obviate the requirement from fairness to run a subsequent lottery when deciding to distribute benefits. This is because fairness requires a lottery with a particular distribution of chances to benefit amongst the outcome groups, not any distribution of chances at all. Different views disagree about which distribution of chances is most fair, but all plausible views do not claim that any distribution of chances are fair. To see this, note that it is not intuitively fair to run a lottery that gives one of the two outcome group in 33v33 Train Cars a 99% chance of benefiting. It is fair to distribute a 50% chance of benefiting to each of the outcome groups in cases where all outcome groups are mutually exclusive and are of equal size, such as 33v33 Train Cars.

In many, perhaps most, cases, all potential beneficiaries will not have had a fair chance to be in the outcome group that is benefited. After all, in *33v33 Train Cars* that would require that each potential beneficiary had a precisely equal chance of benefited. In cases where the prior lottery was unfair, a subsequent

non-natural lottery is required by fairness. So while in a limited set of cases subsequent lotteries will not be required by fairness as the requirements of fairness have already been met, this does not provide a general argument against the fairness of running non-natural lotteries in all cases.

If the prior lottery gave all potential beneficiaries the most fair chance to be in the benefited outcome group, then I agree that a subsequent lottery to distribute chances of benefiting would not be required by fairness. After all, they have already been given a fair chance of benefiting. It does not matter for fairness whether the chances of benefiting distributed are natural or non-natural, just whether or not the chances of benefiting are fair. Is there any argument that for the conclusion that prior lotteries give potential beneficiaries this fair chance in equal conflict cases? Consider the following argument about fairness in equal conflict cases where all outcome groups are equally sized and mutually exclusive:

(P1) Fairness requires that, as far as we know, all potential beneficiaries receive an equal chance of benefiting.
(P2) As far as we know, the prior lottery always ensures all potential beneficiaries receive an equal chance of benefiting.
(C) Therefore, the requirement of fairness in P1 will always be met by the prior lottery.

If (C) is true, then fairness will never require us to run a subsequent lottery as the requirement of fairness has already been met by the prior lottery. However I will argue that while the argument is valid, it is not sound.

It will sometimes, but not always, be true that according to our evidence all potential beneficiaries will have an equal chance of benefiting. If as far as we know (i) all potential beneficiaries could have been in different outcome groups (i.e. that there was a prior lottery) and (ii) have no epistemic evidence that that any potential beneficiary had a greater chance of being in any outcome group than any other, then as far as we know all potential beneficiaries had an equal chance of benefiting. If we only know (ii), then it is possible, according to our evidence, that all potential beneficiaries had no chance of being in a different outcome group and so there was no 'prior natural lottery'. If we only know (i), as far as we know all potential beneficiaries could have had unequal chance of being in the benefited outcome group. Not all cases have these two epistemic features (i) and (ii). For example, we may have good evidence that one beneficiary was more likely to be in one outcome group than another. For example, in *33v33 Train Cars*, if we knew that someone only ever rides in a certain train car due to habit, we have evidence that not all potential beneficiaries have an equal (in this case, fair) chance of being in either outcome group. As a result, P2 is false. As far as we know, the prior lottery does not *always* give all potential beneficiaries an equal chance of benefiting. The argument is unsound. The argument could be modified in an attempt to make it sound.

One could weaken P2 and C, replacing them with P2* 'as far as we know the prior lottery sometimes gives all potential beneficiaries an equal chance of benefiting' and C* 'the requirement of fairness in P1 is sometimes met by the prior lottery'. It certainly is possible for us to have evidence that all potential beneficiaries have an equal chance of benefiting in some equal conflict cases, so I see no reason to reject P2*. However, there is a more fundamental problem with the argument. Even if as far as we know, the prior lottery gives each potential beneficiary an equal, fair chance of benefiting, it does not mean that the prior lottery *actually* gives potential beneficiaries an equal, fair chance of benefiting[58]. To see this, imagine a case in which thousands buy tickets for an equal chance to win money in the next draw of a state run jackpot lottery. In this lottery, money that has not previously won accumulates and is distributed to later lottery winners. However, the machine that distributes chances of winning is temporarily faulty, unbeknownst to everyone, so that it will never select any winner in the next draw of the lottery. The machine works correctly at other times. Even if as far as everyone knows they have a fair chance of winning in the next lottery, the lottery clearly is not fair. It does not give the lottery ticket buyers the actual chance of benefiting in the next draw of the lottery that they have a claim on. Therefore P1, like P2, is also untrue. The requirements of fairness are not met merely because as far as we know all

[58] In the philosophy of probability, there are numerous conceptions of probability, including the epistemic conception in use here. For the purposes of this argument, it is not necessary to endorse a particular view about which conception of probability is most philosophically plausible. The argument in this section of the paper is only intended to show that even if as far as we know all potential beneficiaries had an equal chance of benefiting in the prior lottery, the prior lottery can be unfair and thus fairness may still require us to run a subsequent lottery.

potential beneficiaries have an equal chance of benefiting. I can see no general reason why prior natural lotteries will always meet the fairness of requirements, even in equal conflict cases where all outcome groups are equally sized and mutually exclusive. As a result, in some, perhaps most cases, fairness will require a non-natural lottery to be run to select whom to benefit.

I have argued that a prior lottery only obviates the requirement from fairness to run a subsequent lottery when the prior lottery distributes fair chances of benefiting. However, there is no general reason to believe the prior lottery distributes such fair chances. Whether or not the prior lottery distributes fair chances is a contingent matter and therefore in many cases we will still have reasons from fairness to run a fair subsequent lottery. With this in mind, we can answer Harsanyi's second rhetorical question. The reason why some non-natural lotteries, like those run by governments, are morally superior (with regards to their fairness) to natural lotteries are because prior lotteries often do not distribute fair chances of benefiting, whereas subsequent non-natural lotteries can do so.

VII.C - Unequal Benefits

So far I have only advocated lotteries in equal conflict cases where, by definition, all potential beneficiaries can receive an equally sized benefit. In practice, however, it is extremely unlikely that such cases would arise and if lotteries

were only appropriate for when precisely equal benefits were being distributed, their significance would be extremely limited. What about cases when the benefits being distributed are unequal?

When the potential benefits to equally deserving are unequal, I argue that fairness still requires us to use a lottery to distribute the benefit. There is a particularly compelling example from Broome supporting this view. Broome asks us to consider how we use lottery procedures to decide which player in a game starts in a favourable position. We flip a coin to see who plays white in chess (or black in Go), flip a racquet to see who serves first in tennis and throw dice to determine who goes first in Monopoly. We do this to ensure fairness; however, as Broome as noted, we still think holding such a lottery is fair even when the benefit distributed is unequal. For example, some players receive more benefit than others from being able to serve first in tennis or make the first move in chess. We hold this belief about fairness even in cases where the stakes are high, such as in World Cup tennis tournaments. Furthermore, note that these lotteries are fair even when the benefits to other non-claimants vary. Some players are more popular than others, so we can sometimes predict that if popular players win, there will be more overall happiness than if they lost because of the joy their supporters will experience if they win. We can therefore predict that more happiness is likely if we give the popular player a favourable start. The supporters have no claim on having their supported player have this favourable start, and it is intuitively still fair to run the lottery.

Conclusion

When you cannot benefit everyone in an equal conflict case, fairness requires that you give claimants an appropriate chance at benefiting. In Section II I argued that unweighted lotteries were unfair and implausible because they are insensitive to additional peoples' claims when they are added to existing outcome groups, indeterminate in *unequal chance conflict cases* and particularly unfair in *equal zero benefit cases*. In fact, because unweighted lotteries are indeterminate in unequal chance conflict cases, any lottery theory that purports to account for all equal conflict cases must be a weighted lottery theory. All of these problems arose because unweighted lotteries do not take into account the differing sizes and compositions of outcome groups. In Section III I evaluated all of the weighted lottery accounts of how we should distribute benefits in equal conflict cases currently available in the literature and presented pressing objections to all of them. None of the pre-existing weighted lotteries tell us how we should distribute benefits when claimants can be benefited in more than one outcome group. For these reasons, it turns out that the debate between weighted and unweighted lottery theorists has been misguided and in fact the pressing task for lottery theorists is to determine which weighted composition sensitive lottery approach to adopt. I then considered possible variations of current weighted lottery approaches in Section IV, such as a Timmermann inspired iterated individualist lottery, and argued in Section V that fairness requires us to adopt my exclusive composition sensitive lottery. This exclusive lottery can both be independently motivated

and accords with moral judgments about *convergent cases*. In Section VI I argued that even though fairness gives us reasons to run my exclusive composition sensitive lottery, the fair procedure is not always the morally right procedure. In Section VII I defended lottery distribution views such as my own from three important objections - that weighted lotteries are under motivated; that prior 'natural' lotteries obviate the requirement to run subsequent lotteries and that lotteries have limited application and importance if they are only fair when precisely equal benefits are being distributed.

If I am right about the fair distribution of benefits in equal conflict cases, then there are many implications for both public policy and private moral behaviour whenever such a case arises. It is therefore important to discover which real life distribution problems are equal conflict cases and then to determine if it is feasible to implement a fair exclusive lottery at an acceptable cost. I will not discuss either issue in depth in this paper but I will mention one case of medical resource distribution. The healthcare distribution of donor lungs strikes me as relevantly similar to the cases discussed in this paper. The potential lung recipients considered in such cases are sometimes citizens that have defeasible and equal claims on their healthcare system to be treated. However, some recipients need one lung (e.g. patients with end stage lung disease), whereas others need two (e.g. cystic fibrosis). Given there are far fewer donor lungs of this kind available than potential recipients and the benefit is roughly equal, I contend that fairness requires that an exclusive composition sensitive lottery to be run between potential beneficiaries. This would be one step of many in

making the world a fairer place – which other steps are worth taking is a matter for future research.

Bibliography

Broome, John. *Ethics Out of Economics.* Cambridge: Cambridge University Press, 1999.

Goodwin, Barbara. *Justice by Lottery.* 2nd ed. Exeter: Imprint Academic, 2005.

Harsanyi, John. "Nonlinear Social Welfare Functions: Do Welfare Economists have a Special Exemption from Bayesian Rationality?" *Theory and Decision* 6, no. 3 (1975): 311-332.

Hirose, Iwao. "Review Article: Aggregation and Non-Utilitarian Moral Theories." *Journal of Moral Philosophy* 4, no. 2 (2007): 273-284.

Hirose, Iwao. "Weighted Lotteries in Life and Death Cases." *Ratio* 20, no. 1 (2007): 45-56.

Hooker, Brad. "Fairness." *Ethical theory and moral practice* 8, no. 4 (2005): 329-352.

Kamm, Frances Myra. *Intricate Ethics : Rights, Responsibilities and Permissible Harm.* New York, N.Y. ; Oxford: Oxford University Press, 2007.

Kamm, Frances Myra. *Morality, Mortality: Volume 1.* New York ; Oxford: Oxford University Press, 1993.

Katz, Leo. *Bad Acts and Guilty Minds : Conundrums of the Criminal Law.* Chicago ; London: University of Chicago Press, 1987.

Knight, Carl, and Zofia Stemplowska. *Responsibility and Distributive Justice.* Oxford University Press, 2011.

Lang, Gerald. "Fairness in Life and Death Cases." *Erkenntnis* 62, no. 3 (2005): 321-351.

Mathieu, D. *Organ Substitution Technology: Ethical, Legal, and Public Policy Issues.* Westview Pr, 1988.

Otsuka, Michael. "The Fairness of Equal Chances." (Unpublished Manuscript).

Otsuka, Michael. "Saving Lives, Moral Theory, and the Claims of Individuals." *Philosophy & Public Affairs* 34, no. 2 (2006): 109-135.

Peterson, Martin. "The Mixed Solution to the Number Problem." *Journal of Moral Philosophy* 6, no. 2 (2009): 166-177.

Scanlon, Thomas. *What we Owe to each Other.* Cambridge, Mass. ; London: Belknap Press of Harvard University Press, 1998.

Smart, John, and Bernard Williams. *Utilitarianism; for and Against.* Cambridge University Press, 1973.

Stone, Peter. "On Fair Lotteries." *Social Theory and Practice* 34, no. 4 (2008): 573-590.

Taurek, John. "Should the Numbers Count?" *Philosophy and Public Affairs* 6, no. 4 (1977): 293-316.

Can't Get No Satisfaction - The Chancy Satisfaction of Moral Claims[59]

Many moral philosophers and policy makers contend that when we have to distribute scarce, indivisible benefits amongst potential beneficiaries, each with an equally strong moral claim on the benefit, fairness sometimes requires us to run a lottery to decide who to benefit[60]. Doing so purportedly satisfies their claims[61]. While I have defended a substantive view elsewhere[62], in this paper I do not wish to defend a particular view about which cases we should run lotteries on, nor what kind of lottery fairness requires, but want to focus on the theoretical question of how such lotteries satisfy beneficiaries' claims, in particular by reference to temporally extended cases. I will argue that temporally extended cases have significant importance for both the plausibility of particular philosophical theories about lotteries and the application of benefit-distributing lotteries in practice. Perhaps because this importance is often overlooked, there has been little discussion of the temporal structure of how lotteries satisfy beneficiaries' claims. One exception to this is in John Broome's work, which I will discuss and criticise in Section I. These criticisms illustrate what is needed for a philosophically adequate account of claim

[59] Thanks to John Broome, David Faraci, Johann Frick, Dan Halliday, Julian Savulescu, Jack Woods and audiences at Bowling Green State University, Fordham University, the Princeton Philosophical Society & Oxford University Applied Ethics Discussion Group for comments on earlier versions of this paper.

[60] For example, see Michael Otsuka, "The Fairness of Equal Chances," (Unpublished Manuscript); Gerald Lang, "Fairness in life and Death Cases," *Erkenntnis* 62, no. 3 (2005): 321-351.; Martin Peterson, "The Mixed Solution to the Number Problem," *Journal of Moral Philosophy* 6, no. 2 (2009): 166-177.

[61] John Broome, "Fairness", in John Broome, *Ethics Out of Economics,* (Cambridge: Cambridge University Press, 1999), 111-122.

[62] For my substantive lottery view, see my paper 'Weighing up Weighted Lotteries'.

satisfaction by lottery. Section II of this paper will offer some new alternatives to Broome's view that attempt to meet the theoretical requirements specified in Section I. In this section, I advocate a new view called the *dual structure view*, so called because it posits that claimants have two relevant claims in distributional cases of the above kind. The last part of Section II defends the *dual structure view* against an objection. Section III draws out a practical implications of my view for a variety of temporally extended cases, including the distribution of corneas to patients who have suffered corneal degeneration.

Section I - The Role of Lotteries in Fairness & Broome's Synchronic View

I.A - The Role of Chances of Benefiting for Fairness

Before outlining and objecting to Broome's view, it is important to illustrate the role of *ex ante* chances of benefiting for fairness, for claim satisfaction and for permissibility. In order to illustrate this, consider the following case;

> *The Chancy Doctor.* A doctor can conduct lifesaving surgery on two equally badly off patients in two separate operations. These two patients are the only people she can currently help. The operation can be conducted easily with guaranteed success. Furthermore, she has no other obligations or better things to do with her time. Instead of just conducting the surgery on both patients, she decides to roll two separate six sided dice. If and only if the first die rolls a 1, she will choose to conduct surgery on the first patient. If and only if the second die rolls between 1-5, she will choose to conduct

surgery on the second patient. This is the only difference in the doctor's treatment of the two patients. The first die rolls 1 and the second 4, so she actually conducts successful surgeries on both patients. She does not tell the patients about the lottery she conducted nor does anyone ever find out what she did.

Even though the *Chancy Doctor* gives the same lifesaving benefit she owes to both patients, her action is patently unfair and wrong because of the differences in the *ex ante* chances of benefiting she gives to the two patients. It would be even less fair if she gave one patient a 100% chance to be treated but gave no chance to the other. It is clear, then, that people's chances of benefiting are relevant to what is fair and what is wrong. While I will not assume any contentious theory about *which* chances of benefiting are fair in all cases, for the purposes of this paper I will assume that *some* distributions of chances are fair. The Chancy Doctor does not appropriately satisfy the potential beneficiaries' claims. By claims, I mean a particular kind of reason, namely acts that are owed to a particular individual or group. Claims are patient-relative reasons, where a patient is a person or thing that undergoes an action of another (this meaning of 'patient' contrasts with the meaning of 'agent')[63]. Because the doctor can save both of their lives easily with guaranteed success,

[63] Sometimes patients have claims on no particular agent, for example when a claimant could have their premature death easily prevented by any one of a number of potential agents.

fairness intuitively requires her to give both patients an equal 100% chance of benefiting. To do so is the appropriate way to satisfy the patients' claims[64].

I.B - Broome's Synchronic View of Claim Satisfaction by Lottery

How do chances to benefit satisfy potential beneficiaries' claims on scarce, indivisible benefits? According to John Broome, fair lotteries provide surrogate satisfaction of claims by giving all claimants an appropriate chance of receiving the benefit as a surrogate for actually receiving the benefit. His view is synchronic. When such a lottery is run, those that are not selected to benefit by the lottery have received surrogate satisfaction of their claim and thus no longer have any claim on the benefit[65]. Those who are selected by the lottery (i.e. those who win) have their complete claim intact and thus we ought to give the benefit to the potential beneficiaries who still have claims on the benefit.

Broome motivates this view by arguing that it would be absurd to ignore the first lottery and run a second lottery amongst precisely the same potential beneficiaries. Broome's view explains this absurdity. On Broome's synchronic view, after the lottery has been run, the number of available benefits is equal to the number of claimants and so all claimants can be benefited. If there were no

[64] For a more full account of how peoples' claims can be satisfied by chances of benefiting, see John Broome, "Fairness", in ibid.. With the exceptions of the criticisms in this paper, I accept Broome's general view about how lotteries are relevant to fairness when distributing scarce, indivisible benefits. I, however, disagree about what chances of benefiting fairness requires. In 'Weighing up Weighted Lotteries', I argue that fairness sometimes requires giving potential beneficiaries an unequal chance of benefiting.

[65] John Broome, "Uncertainty and Fairness," *The Economic Journal* 94, no. 375 (1984): 624-632.

changes to potential beneficiaries' claims after the first lottery (and no other morally relevant changes), Broome contends, there would be equally good reason to run a second lottery as to run the first. In fact, if there were no changes to the claims between lotteries then any number of iterative lotteries could be just as justified as the first lottery as long as one was able to save the lotteries' winners. Call the implausibility of these iterative lotteries, *iterative absurdity*. According to Broome, it is absurd to tamper with a previously run lottery that has already provided surrogate satisfaction of some beneficiaries' claims by running a subsequent lottery between the exact same potential beneficiaries. This is because, Broome argues, doing so gives beneficiaries with no claims on receiving the benefit a chance of receiving the benefit. While I agree that running subsequent lotteries in such cases is absurd, I will argue that *iterative absurdity* can be avoided without adopting Broome's synchronic view and furthermore that Broome's view has implausible implications.

I.C - Objections to Broome's Synchronic View

While initially plausible, Broome's synchronic view commits him to take morally irrelevant features of some benefit distribution cases as morally relevant and undermines his own view about the fairness of giving claimants an equal chance of benefiting. On Broome's view, lottery winners retain their complete claim on the benefit, whereas lottery losers have no claim on the benefit. This leads to particularly implausible implications in temporally extended cases such as the following:

Stormy seas. On a ship, two people (A & B) fall overboard during stormy weather and will drown if not saved. Each person has an equally strong claim to be saved by the ship's captain. Unfortunately those overboard can only be saved through the use of life buoys and there is only one such buoy which can be used to save just one person. Call each set of people (sometimes a singleton) that can be benefited together an *outcome group*. In this case, A is an outcome group and B is another outcome group[66]. There are no morally relevant differences between either outcome group. Having read Broome[67], the captain opts to run a lottery that gives each claimant an equal chance (50%) of winning the lottery and selects one winner (B). Given his long experience with games of chance, he is able to run such a lottery quickly with ease and without endangering the lives of those he is about to save. However, after he has run the lottery but before he throws out the buoy, another wave hits the boat and two more sailors (C & D) fall overboard. What should the captain do now?

[66] In order to be ecunemical between different views on what the appropriate chance of benefiting in such cases is, I will assume that all outcome groups are of the same size, do not have overlapping members and that all of the benefits are equal in all morally relevant respects. In such cases, all lottery approaches extant in the philosophical literature claim that all potential beneficiaries ought to have an equal chance of benefiting (or that it is fair that they have an equal chance of benefiting). For more on this, see the discussion of weighted and unweighted lotteries in Section II.

[67] Actually, as this is one of the least contentious cases in which we ought to run a lottery, most writers that advocate lotteries in some cases advocate running a lottery in this case, including Hirose, Timmermann, Rakowski & Saunders. See Iwao Hirose, "Weighted Lotteries in Life and Death Cases," *Ratio* 20, no. 1 (2007): 45-56.; Jens Timmermann, "The individualist lottery: how people count, but not their numbers," *Analysis* 64, no. 2 (2004): 106-112.; Ben Saunders, "A defence of weighted lotteries in life saving cases," *Ethical theory and moral practice* 12, no. 3 (2009): 279-290.; Eric Rakowski, "Taking and saving lives," *Columbia law review* 93, no. 5 (1993): 1063-1156.

On Broome's synchronic view, because A has lost the lottery, A has no claim on receiving the benefit. However, B, C & D all have equally strong claims on the benefit, but not all of them can be benefited, so according to Broome's synchronic theory, fairness requires the captain to run a second lottery between them which gives them all equal chances of winning (1/3rd). According to Broome, when one has to distribute a scarce, indivisible benefit amongst potential beneficiaries who all have an equally strong claim on the benefit, fairness requires us to run a lottery that gives each claimant an equal chance of winning the lottery[68]. While doing so gives B, C & D equal chances of winning the second lottery, it does not give all potential beneficiaries (A through D) equal chances of receiving the benefit. Those who fall overboard later (in this case, C & D) have a greater chance of receiving the benefit as they participate in fewer lotteries and thus have less chance of losing their claim. In this case, C & D have to win just one lottery which gives them 1/3rd chance each in order to be saved. However those that fell in the beginning (A & B) have to win two lotteries, one with 50% chance of winning, one with a 1/3rd chance of winning, in order to be saved. As such, they have a significantly lower chance of being saved (1/6th) than those that fell overboard later. This is implausible as it gives the time the people fell overboard excessive moral importance.

This is more implausible the more subsequent groups fall overboard. Call potential beneficiaries that have their claims on the benefit created after a lottery has been run *subsequent claimants*. If a third group of *subsequent*

[68] Also, as the outcome groups are of the same size and mutually exclusive, all lottery theorists I know of would give all claimants an equal chance of winning the lottery.

claimants fall overboard after the second lottery has been run, then there is more discrepancy between potential beneficiaries' chances of benefiting. On Broome's synchronic view, the later a group develops its claim on the benefit (and thereby develop their claim after one or more lotteries have been run), the greater chance they have to benefit. This implausibly gives an intuitively morally irrelevant feature of this case great moral importance. As it is a structural feature of the case, not the particular details, that is necessary for the objection, the objection can be generalized. Furthermore, this is not merely a philosopher's example in a distant possible world, but there are structurally similar actual world cases. For example, if lotteries run in accordance with Broome's synchronic view were used to determine who should receive donor kidneys, those that suffered from renal failure more recently would have a greater chance of receiving a kidney.

As it is the time that people develop their claims relative to the time the lotteries are conducted that affects their chances of receiving the benefit, Broome's synchronic view also has implausible implications in another kind of temporal case. Imagine *Stormy Seas* except that because the captain was slightly slower in running his lottery, that C, but not D falls overboard before he can run the lottery. This changes the case so that D has a 50% chance of being saved but A through C only each have $1/6^{th}$ chance of being saved. Again, Broome's view implies that a feature of the case that is morally irrelevant (the slight difference in speed in which one conducts a lottery) has been elevated to something of life and death importance.

Another type of temporally extended case also presents problems for Broome's account of claim satisfaction. Imagine *Stormy Seas*, except that instead of C & D falling in after the first lottery has been run, the captain finds a second life buoy, equally as good as the first, with which he can save another person with. Call this case *Extra Buoy*. It seems absurd in *Extra Buoy* to claim that while B has a claim to be saved as he won the first lottery, A has no such claim. Broome has suggested to me in conversation that this apparent absurdity can be explained away by considering specifically what A & B have a claim on. If, prior to the lottery, they have a claim not just on the general benefit of having their life saved, but on the specific benefit of having their life saved by the first buoy, then this explains why the lottery loser, A, has no claim on the first buoy. Call this the *specific benefits view*. However, it is not plausible that A & B have a specific claim of the sort Broome has suggested. To see the implausibility of this, consider *Extra Buoy* again. According to the specific claim view, B could complain if they are given the second buoy instead of the first as it would not satisfy their claim. Insofar as it is better to satisfy moral claims than not, all other things being equal, according to the specific benefits view, it is better to give the first buoy to B and the second buoy to A as opposed to giving the first buoy to A and the second buoy to B because it satisfies B's specific claim on the first buoy. Intuitively, it does not seem to matter whom gets which buoy, as long as both of their lives are saved. Again, Broome's view raises something that is intuitively morally irrelevant to something of significant moral importance.

Another problem for Broome's theory emerges if we try to extend it in plausible ways to cases where the lotteries run are not perfectly fair. Consider a case where we can just save one of two lives, as in the start of the *Stormy Seas* case. Now stipulate that in order to decide who to benefit, a slightly unfair lottery is run - one that gives A a 49% chance of benefiting and B a 51% chance of benefiting[69]. It is unfair because equally deserving claimants that are not morally different in any way are treated (slightly) unequally. Broome does not consider this case, so I have to consider how his theory could be extended in a consistent way to account for it. As the lottery is not completely fair, as it would be if they were each given a 50% chance of benefiting, it is implausible that the lottery would result in either claimant having their claim completely satisfied or having no claim on the good after the lottery. What, then, is the most plausible

[69] In many actual cases of distribution, precise control over the chances of benefiting cannot be obtained or is not feasible. This is sometimes due to inability to run precise enough lotteries (e.g. obtaining perfectly fair coins and a perfectly fair coin toss) or features about the way in which beneficiaries are organised into outcome groups. To take an example from my college - spots at special dinners are scarce, as there are more people that want to attend than are able to be seated in my college's hall. Furthermore, every college member who wants to attend has an equally strong claim on the benefit of attending the dinner. However, as is often the case, they would not benefit from the dinner if they were not able to go with their college friends on the same table, the college's hall is split up into a number of tables and the college is large enough that not everyone in the college are friends with each other. So some members can only be benefited in certain outcome groups. As it would not be feasible to run a lottery between all possible ways in which outcome groups could be seated on the various tables, it is reasonable to run a lottery that compromises feasibility and fairness, by running an iterative lottery that selects one outcome group after another until all places are filled. However this kind of lottery will give greater chances of benefiting to smaller outcome groups (e.g. couples), than larger ones, as the lottery reaches the end of each table, when there are likely to be less spots available, thereby excluding larger outcome groups from being included in the lottery. To illustrate, imagine we have a table with 8 spots, and the outcome groups are as follows A&B&C&D&E&F, G&H&I&J, K&L&M&N&O, P&Q&R&S, T&U, V&Y and Z. Say, for example, that the group of five is selected (K&L&M&N&O), then all groups containing four or more members cannot be included in the subsequent lottery. Structurally analogous cases would occur whenever the number of spots left are smaller than some so far un-benefited outcome group. This would give smaller groups an objectively higher, unfair, chance of benefiting due to feasibility concerns and the way in which outcome groups are structured. As those who have organised weddings will know, distributing seats on wedding dinner tables pose some of the same distributional problems, but even more complexities arise!

way of extending Broome's theory to account for this case, where they are treated in a way that is almost fair according to Broome? It would be to treat them as if they have had their claims partially satisfied[70]. That is, the claimants each have less of a claim on the good after the lottery than before, but they still have some (weakened) claim on the good. How much are those claims partially satisfied?

Here, a Broomean theory faces a dilemma. On one horn would be the view that claimants would have their claims *equally* partially satisfied, as they are given chances of benefiting equally distant (1%) from the fair chance of benefiting (50%). That would result in each claimant having a claim on the good that is less strong than they had previously but equally as strong as each other's claim. There are two problems with this horn. Firstly, if both beneficiaries still have claims that are equally strong, then the objectionable *iterative absurdity* arises. In fact, this *iterative absurdity* is even more implausible than the original *iterative absurdity*, as each time an imperfect lottery is run, their claims are weakened, until they have nearly no claims at all. Secondly, that they each have their claim equally partially satisfied is implausible. Consider the following analogy. If you owe two people each $1 and give one 51c and the other 49c, you have not equally satisfied their claims. As such, it seems that in this imperfectly fair case, that their claims are *unequally* partially satisfied, which is the other horn of the dilemma. On this view, in line with Broome's view that lottery losers in fair lotteries completely lose their claim, the loser of the somewhat unfair

[70] It is, after all, more fair than running a lottery that just gives A a 100% of benefiting and B no chance of benefiting, and each claimant receives *almost* a fair chance of benefiting.

lottery has a less strong claim on the benefit than before. In order to avoid *iterative absurdity*, this claim will have to be weakened enough so that it is no longer fair to run a lottery between previous lottery losers and winners. But such a view will give rise to the same implausible implications we saw before in temporally extended cases like *Stormy Seas*. It will claim that previous lottery losers have insufficiently strong claims to be included in subsequent lotteries, so *subsequent claimants* will have greater chances of benefiting than claimants that developed their claims prior to a lottery being run. As Broome's theory already has problems with *subsequent claimants* and is motivated by its avoidance of *iterative absurdity*, it would be best for a Broomean to embrace that horn of the dilemma.

These cases show that upon closer inspection, Broome's synchronic view of the way in which lotteries satisfy people's claims do not directly give claimants equal chances to receive the benefit, but instead give claimants equal chances to *lose* their claims in each lottery. Those that win the lottery are not any better off (in an absolute sense) than they were before the lottery as they have the same claim both before and after the lottery. Broome's lotteries only indirectly determine potential beneficiaries' chances of receiving the benefit as those that receive the benefit are those that still have their claims after the last lottery that according to Broome results in the number of claimants equaling the number of available benefits. This feature not only causes Broome to elevate morally irrelevant features of these cases to something of great importance, it undermines one of his other views. Broome contends that fairness in cases

such as these requires all claimants to get an equal chance of receiving the benefit. Broome's synchronic view in temporally extended cases such as *Stormy Seas* give each claimant an equal chance of winning each lottery but different chances of actually receiving the benefit. On Broome's account, it is equal chances of receiving the benefit, not equal chances of winning each lottery, which are the fair surrogate for actually receiving the benefit. Broome's theory explains the unfairness of the second synchronic lottery in *Stormy Seas* well – it gives A & B each a $1/6^{th}$ chance of surviving compared to C & D's $1/3^{rd}$ chance of survival when fairness requires them to have an equal 25% chance of surviving. As a result, Broome's synchronic view of claim satisfaction by lottery undermines his view of fairness in temporally extended cases such as *Stormy Seas*. This is because it results in running lotteries that are unfair by Broome's own theory and the reason for running lotteries in the first place was to ensure a claimant selection procedure that was as fair as possible. I will consider and reject a possible response to this undermining objection in the next section.

While Broome's synchronic theory is able to explain *iterative absurdity*, it does not give a plausible account of temporally extended cases where either *subsequent claimants* appear (e.g. *Stormy Seas*) or additional benefits become available (e.g. *Extra Buoy*), or cases where only partially fair lotteries are run. We need a view that plausibly accounts for these four features, and I turn to some alternatives to Broome's view in Section II.

Section II - Alternative Views of Claim Satisfaction

II.A - Diachronic Alternatives

Two diachronic views of claim satisfaction by lottery are more plausible than Broome's in two ways. On one such view, those that lose the lottery have their claim qualitatively weakened but not negated[71]. Call this view *diachronic weakening.* On the other view, those that win the lottery have their claim qualitatively strengthened. Call this view *diachronic strengthening.* Both of these views can explain the *iterative absurdity* of a second lottery between the same claimants of the first lottery. This is because it would be absurd to run a lottery between those with (sufficiently) stronger claims and those with weaker claims[72]. Both views would also explain why both A & B have claims to be saved in *Extra Buoy* (though they claim, somewhat oddly, that A has a less strong reason than B to be saved). Furthermore, these two views are intuitively plausible as they maintain that lotteries that give claimants a less than guaranteed chance of receiving the benefit only partially satisfy their claim on the benefit. After all, intuitively, full satisfaction of a claim on a benefit is only

[71] Note the similarity to the second horn of the dilemma presented against Broome in Section I.

[72] Precisely how much stronger a claim has to be in order to make running a lottery between them absurd and impermissible is a matter for contentious debate. In fact, as Parfit has suggested is true of some reasons, there may be no truth about some claims' relative weight/strength (see Derek Parfit, *On What Matters* Oxford University Press, 2011, Volume 1, 130-141). However, to illustrate the absurdity of running a lottery with a clear cut case, consider the following. We only have one dose of a medicine which can be used to treat one of two people. The medicine can be used to prevent the premature death of one person or stop the mild headache of the other person. Assuming that, in this case, claims depends on the size of the benefit they will receive, the person who will die if untreated will have a much stronger claim than the person suffering a headache. It is absurd to think one should run a lottery to decide whom to benefit in this case, one ought to just benefit the person who will die if untreated.

achieved by actually receiving the benefit[73]. This is why it would be fairest to give all claimants the benefit they have claims on, which is possible in cases like *Extra Buoy* and the *Chancy Doctor* but is impossible in cases of scarcity like *Stormy Seas*. In these two ways, both diachronic views are more plausible than Broome's.

The differences between these three views can be illustrated by the following table. The numbers indicate the strength of that individual's claim on the benefit. The starting strength of the qualitatively equal claims is arbitrary (50). However, 0 indicates that the beneficiary has no claim on the benefit. Furthermore, stipulate that when claimants are 10 or more units apart from one another, it is not fair to run a lottery between them to decide who to benefit as their claim is significantly stronger than the other's – this avoids *iterative absurdity*. These numbers are provided only for illustration - it is not essential for my argument that the strength of claims can be precisely measured quantitatively.

[73] Some philosophers, such as Rahul Kumar, claim that lotteries provide partial satisfaction of potential beneficiaries' claims on the good without a complete account of how such claims are partially satisfied. These diachronic views provide a more detailed (but as we will see, ultimately implausible) explanatory account of the *partial* satisfaction of claims. Rahul Kumar, *Consensualism in Principle: On the Foundations of Non-consequentialist Moral Reasoning*, London: Routledge, 2001), 130-132.

	Before First Stormy Seas Lottery		After First Stormy Seas Lottery				Extra Buoy After First Lottery	
	A	B	A	B	C	D	A	B
Broome's Synchronic View	50	50	0	50	50	50	0	50
Diachronic Strengthening	50	50	50	60	50	50	50	60
Diachronic Weakening	50	50	40	50	50	50	40	50

While more plausible than Broome's synchronic view, these diachronic views should also be rejected. This is because both views also lead to the implausible conclusions of Broome's account in temporally extended cases like *Stormy Seas*. *Diachronic weakening* would claim that those who had lost previous lotteries would not have sufficiently strong claims to be included in subsequent lotteries such as *Stormy Seas*. After all, this is how *diachronic weakening* avoids *iterative absurdity*. In this way, *diachronic weakening* is similar to Broome's synchronic view that in temporally extended cases like *Stormy Seas* fairness requires us to run lotteries between previous winners and *subsequent claimants* while excluding previous lottery losers. *Diachronic strengthening* has an even more implausible result. It claims that those that win the first lottery have significantly stronger claims than lottery losers *and* those that fall overboard after the first lottery. On such a view, both lottery losers and *subsequent claimants* have qualitatively equal claims that are insufficient to justify running subsequent lotteries to give them any chance of surviving. As such, *diachronic strengthening* is less plausible than *diachronic weakening*.

What lottery theorists need is a view that avoids *iterative absurdity* while giving a plausible account of temporally extended cases such as *Stormy Seas* and *Extra Buoy*. I will argue that such a view is available once we analyse the claims that potential beneficiaries have more carefully.

There is one theoretically relevant problem with all of the accounts of how lotteries satisfy potential beneficiaries' claims so far discussed which gives us some guidance to a better theory. In order to avoid *iterative absurdity*, they maintain that potential beneficiaries' claims on the benefit change depending on whether or not they win the first lottery. However, then when new *subsequent claimants* are created in a temporally extended situation like *Stormy Seas*, the differing strengths in claims on the benefit that were adopted to avoid *iterative absurdity* lead to implausible implications for lotteries involving *subsequent claimants*. I cannot see how when one focuses solely on claims on the benefit one can solve this problem. Fortunately, this problem can be solved when one recognizes that potential beneficiaries have two types of claims that are relevant to benefit distribution cases like these and that these can be used to explain how lotteries affect peoples' claims.

II.B - The Dual Structure View

By stipulation, in the cases I'm discussing, all potential beneficiaries have qualitatively equal strong claims on the benefit. For brevity, call these their *benefit claims*. These claims are satisfied, *ex post*, when the benefits are actually

distributed. In addition to these *benefit claims*, I posit that they also each have a claim not to be treated unfairly - this is a matter of procedural *ex ante* fairness. Call these their *procedural claims*. It is these *procedural claims* that are satisfied by a lottery that gives claimants appropriate chances of benefiting. I call my view that posits both *benefit claims* and *procedural claims* in these cases the *dual structure view*. The dual structure of the claims parallels the dual structure exhibited in the *Chancy Doctor* case - even though both patients were successfully given the life saving treatment they had a claim on and thus had their *benefit* claims fully satisfied, the Doctor's act was wrong and the patients were treated unfairly. This is because while their claim on the benefit was satisfied, their claim to be treated procedurally fairly was not.

Furthermore, my dual structure view also explains *iterative absurdity,* namely why there is no reason to run a second (or a third, or a fourth...) lottery with the exact same claimants. On my dual structure view, the second procedural claim has already been met and running a second lottery would be to invalidate a previously fair situation for no reason. Furthermore, on this view, unlike Broome's, some claims that potential beneficiaries have on the benefit cannot be lost due to results of a lottery. As noted in this paper's title, those benefit claims *can't get no satisfaction*, not even partial satisfaction, through a lottery. This explains why people who lose the lottery still have as strong claims on receiving the benefit and as such offers the best account of temporally extended cases where the number of benefits available for distribution increase, like *Extra Buoy*. In *Extra Buoy*, both A and B both have equally strong benefit

claims after the lottery. The dual structure view does not maintain, like all of the synchronic and diachronic views discussed above, that A is less deserving of being saved. In such a case, according to my view, the earlier lottery can be ignored as it is no longer the most fair distribution procedure. Instead of giving A or B each a 50% chance of benefiting, in *Extra Buoy*, both claimants should receive a 100% chance of benefiting, like in the *Chancy Doctor* case. On my view, while it is better for potential beneficiaries to get some chance at benefiting, preferably the most fair chance, that in itself does not satisfy their *benefit claims*. That can only be achieved when they actually receive the benefit that they have a claim on. When it distributes the most fair chance of benefiting, the lottery procedure satisfies potential beneficiaries' *procedural claims*, not their *benefit claims*.

My dual structure view also explains why it is fair to run a new lottery that gives all potential beneficiaries an appropriate chance of benefiting in temporally extended cases such as *Stormy Seas*. This is because when *subsequent claimants* are added previously conducted lottery procedures are no longer fair. They are no longer fair as on all views of the fairness of lotteries in benefit distribution cases, the fair chances of receiving the benefit depend on the number of outcome groups (and sometimes, furthermore, on the group's size and composition). There are two types of views about the fair way to distribute chances amongst claimants. On unweighted lotteries, all claimants are always given an equal chance of receiving the benefit. On weighted lotteries, claimants are sometimes given an unequal chance of receiving the benefit. Weighted

lotteries give greater chances of receiving the benefit to larger outcome groups – they weight chances by the number of claimants in each outcome group. However, in this case where all outcome groups are of equal size and are mutually exclusive, they both recommend giving all claimants equal chances of receiving the benefit. By adding *subsequent claimants* in cases like *Stormy Seas*, the fair chances of receiving the benefit change and as such the previously conducted lotteries are no longer fair. On both weighted and unweighted lotteries, the fair chance of each potential beneficiary to receive the benefit is 25%. As such, the original lottery that only included A & B is not fair (according to both unweighted and weighted lotteries) when additional claimants C & D are added as in *Stormy Seas*. This explains why procedural fairness requires us to run a new lottery procedure between all potential beneficiaries in *Stormy Seas* and ignore the now unfair results of previous lotteries, just like the 50% lottery in *Extra Buoy* gets ignored once the second buoy becomes available.

In a similar way, the dual structure view can be used to explain lotteries that are not completely fair. When imperfectly fair lotteries are run, like the 49%/51% lottery discussed in Section I.C, the claim not to be treated unfairly is not fully satisfied – it is merely partially satisfied. It would more fully satisfy that procedural claim (and be fairer) to ignore the results of a previously run unfair lottery and run a new, fair lottery.

So the dual structure view provides a philosophically adequate explanation of all of the four desiderata set out in Section I - it explains *iterative absurdity*,

temporally extended cases where either *subsequent claimants* (e.g. *Stormy Seas*) or additional benefits appear (e.g. *Extra Buoy*), as well as cases where only partially fair lotteries are run. Furthermore the dual structure it posits parallels the structure evident not just in the abovementioned cases of benefit scarcity (i.e. when not all potential beneficiaries can be benefited) but also cases like the *Chancy Doctor* where all claimants can be benefited.

II.C - Defending the Dual Structure View

Broome has objected that lottery winners of earlier, previously fair lotteries would have a justified complaint that their earlier winnings are ignored and this complaint would be particularly salient if earlier losers won a subsequent lottery. Broome has suggested in conversation that this may be avoided by running subsequent lotteries that take into account earlier results by only running subsequent lotteries between previous winners and subsequent claimants, while adjusting the probabilities of subsequent lotteries to ensure all claimants have an appropriate chance of benefiting. For an example of how this would work, consider *Stormy Seas* again. Instead of running a subsequent lottery that gives either B, C & D or A, B, C & D an equal chance of winning, one may run a lottery that gives B a 50% of winning, and C or D each a 25% chance of winning. This would ensure that A, a previous loser, cannot be saved by a subsequent lottery, while ensuring all claimants A, B, C and D have an equal (25%) chance of being saved. For A or B to be saved, they each have to win two lotteries each of which they have 50% chance of winning, so their chance of benefiting is

0.5x0.5=0.25. For C or D to be saved, they each have to win only one lottery, of which they have a 25% chance of winning. Note that on Broome's synchronic view, this account will result in giving claimants with claims of equal strengths (*subsequent claimants* and former lottery winners), different chances of winning subsequent lotteries.

However, I will object to Broome's claim that lottery winners of earlier lotteries have a justified complaint that their earlier winnings are ignored. To see that they have no such justified complaint, consider a variation on the *Extra Buoy*[74]. Call it the *Defective Extra Buoy*. The difference is that the first buoy that is available is faulty in that one side of it is more buoyant than the other, so while it can be used to save someone, being saved using it involves significant distress. This is because the defective buoy is not a stable flotation device and results in the person being saved being dunked under the water multiple times before being saved. Now imagine that after the first lottery is run, a normal, non-defective buoy becomes available. In this case, it seems absurd that the first lottery winner would have a complaint if their earlier win on the lottery for the defective buoy was ignored and they were saved by the non-defective buoy. This case, too, illustrates an objection to the *specific benefits view*. It is absurd to think that the first claimant who wins (and thus has his claim intact) has a claim on the defective buoy and should be given that particular defective buoy, and the other potential beneficiary who lost the lottery should be given the non-defective buoy. On such a view, the first lottery winner would have a justified

[74] Thanks to Johann Frick for this suggestion.

complaint if they were given the non-defective buoy instead of the defective buoy to which they purportedly have a claim to be given. This is highly implausible. However if the first lottery winner has no justified complaint in *Defective Extra Buoy*, why should they have a complaint in the case where the potential benefits are equal, as in *Stormy Seas*? I see no principled, non ad-hoc distinction that can be plausibly drawn between these two types of cases and as such reject the claim that lottery winners of earlier lotteries would have a justified complaint that their earlier winnings are ignored. Next, I'll consider a context sensitive view that presents an alternative to my dual structure view while making the same prescriptions as mine in temporally extended cases.

It has been suggested that one could defend a broadly Broomean view that only posits one kind of claim (benefit claims) by making such claims context sensitive[75]. If such a view could account for all of the temporally extended cases discussed earlier, it is supposedly theoretically preferable to my dual structure view as it would be more parsimonious by positing fewer kinds of claims. The kind of view suggested was that one did not have claims on the benefit *simpliciter*, but instead claims on the benefit given a certain claimant context. On such a view, when the number or identity of the claimants changes, so do the claims. Applied to *Stormy Seas*, A&B each have a claim on the benefit when claimants A&B are present, but when C&D fall in, A&B&C&D all each have a *new and different* claim on the benefit when claimants A&B&C&D are present. Their claims are sensitive to the claimant context, so call a view of this kind *claimant*

[75] Such a view was suggested to me independently by Jack Woods and David Faraci in separate discussions.

sensitive contextualism. This view could then be used to explain the same act prescriptions for *Stormy Seas* I proposed in Section II.B. According to this contextualist view, the fact that A lost their claim when A&B were the only claimants present is irrelevant once two new claimants (C&D) have fallen into the water, as all of the claimants have new claims unaffected by the previous lotteries. As such, it explains why given the change in the claimant context, a new lottery that gives all claimants a fair chance to benefit should be run.

However *claimant sensitive contextualism* is unable to explain temporally extended cases where the number of benefits changes but claimants remain the same. This is because as stated so far, it is only sensitive to changes in claimants, not benefits and therefore cannot be used to explain cases like *Extra Buoy* in which only the available benefits change. One could modify the view to be sensitive to both changes in benefits and claimants, and give it a new name, *benefit-claimant sensitive contextualism*. On such a view, when either the benefits or the claimants change, the claimants have their claims changed and lotteries that affected earlier, different claims are irrelevant to new contexts. In order to account for cases in which the available benefits are of different magnitudes, such as *Defective Extra Buoy*, such a context sensitive view would have to be sensitive to both the number and magnitude of the benefits. So in cases like *Extra Buoy,* the fact that A had lost the first lottery would be irrelevant to their claim on the benefits when the second benefit becomes available.

While I agree that such a view could capture temporally extended cases where the benefits available and/or the claimants change, it cannot account for important cases like the *Chancy Doctor* or lotteries that are not completely fair. This is because in these cases, neither the claimants nor the benefits change. As there is no contextually relevant change in the claimants or benefits, *benefit-claimant sensitive contextualism*, cannot be used to explain lotteries that are not completely fair. As such, such a contextualist view would be subject to the dilemma I presented against Broome in Section I.C. Furthermore, according to *benefit-claimant sensitive contextualism*, in the *Chancy Doctor* all claimants receive the benefit that they have a claim on, so a failure to satisfy their benefit claims cannot be used to explain the situation's unfairness. In order to explain cases like this where the context does not change, a contextualist would have to adopt a view similar to my dual structure view. This should not be surprising, as the dual structure of my view about claims parallels the dual structure exhibited in *Chancy Doctor* case. Since we have to adopt a dual structure view to explain the *Chancy Doctor* and imperfectly fair cases *and* such a view can already explain temporally extended cases where the number of claimants and/or benefits, there is no good reason to posit context sensitive claims. The dual structure view explains all of the cases in a more theoretically parsimonious way. Furthermore, even if it was possible to explain the *Chancy Doctor* and imperfectly fair cases without positing a dual structure view, it is not clear that *benefit-claimant sensitive contextualism* is actually a more theoretically parsimonious view than my *dual structure view*. Instead of positing two kinds of claims, it posits one very complex type of claims that are sensitive to the

number and identity of claimants as well as the number and magnitude of benefits. Each time any of these features change, all of the claims change. As such, the view's only purported advantage is contentious, at best.

Not only does my dual structure view meet the four desiderata set out in Section I, it has many important practical implications regarding temporally extended benefit distribution cases. I turn to these in the next section.

Section III - Practical Implications

The dual structure view I have advocated has practical implications for a variety of temporally extended scarce benefit distribution cases. In all of these cases, I will assume that there are no reasons to not give the benefits to the claimants and there are no other ways in which we can utilise those benefits. Furthermore, as before, to ensure weighted and unweighted lotteries make the same lottery prescriptions, I will assume that all outcome groups are mutually exclusive, of the same size and consist of equally deserving claimants[76]. I will first discuss simple temporal cases when either the number of potential beneficiaries increase or decrease or the number of available benefits increase or decrease and no other morally relevant features change. Then I will discuss more complex cases - when both the numbers of potential beneficiaries and

[76] For a discussion of how we ought to balance fairness and the benefits distributed when outcome groups are not the same size and/or contain some of the same members, see my paper 'Weighing up Weighted Lotteries'.

available benefits change; and furthermore when the quality of available benefits change.

The simplest change that can occur in these distributional cases is when the number of available benefits increases so that all claimants can be benefited with the benefit they have a claim to - that is, when there is no scarcity of the benefits that claimants have a claim to. In such cases, all other things being equal, we ought to give all claimants a 100% chance of benefiting - doing so satisfies both their claim on the benefit and their claim not to be treated unfairly. This matches with our moral judgments regarding what ought to be done in the *Chancy Doctor* and *Extra Buoy* cases. Similarly, when both the numbers of potential beneficiaries and available benefits change so that all claimants can be benefited with the benefit they have a claim to, all other things being equal, then we ought to give all claimants a 100% chance of benefiting.

When the number of available benefits increase but one still cannot benefit all claimants, then the fair thing to do depends on whether the initially available benefits have already been distributed or not. Assuming the initially available benefits have already been distributed fairly, then one can treat the *remaining claimants* as a new case of scarce benefit distribution. *Remaining claimants* are those that previously received a fair chance of benefiting but did not actually benefit when the benefits were actually distributed. Fairness requires one to run an appropriate lottery between *remaining claimants* when new benefits become available. When all outcome groups are of equal size, all possible

distributions of the lottery will benefit the same number of claimants and satisfy an equal number of *benefit claims*. Given scarcity, running a lottery to decide whom to benefit would best satisfy their dual claims. If initially available benefits have not yet been distributed, then the fair thing is to run a new lottery between all claimants. This is because on all of the lottery approaches in the literature (both weighted and unweighted), the initial lottery is no longer fair, as the fair lottery depends on who can be benefited and the number of benefits available. This is similar to what fairness requires when the number of claimants increases (i.e. cases with *subsequent claimants*) but the benefits are yet to be distributed.

The addition of *subsequent claimants* when the benefits are yet to be distributed changes what chances of benefiting are fair and so fairness requires that a new lottery be run in order to best satisfy beneficiaries' *procedural claims*. When there are *subsequent claimants* but all available benefits have already been distributed, there is not much one can do as one cannot give claimants any benefits. This is especially bad for the *subsequent claimants* - at least the remaining claimants that were not selected to be benefited previously received a fair chance of benefiting. *Subsequent claimants* in cases where all the benefits have already been distributed never receive any chance of benefiting, unlike the equally deserving claimants before them. As such, fairness requires us, all other things being equal, to wait until all claimants have developed their claims before distributing scarce benefits. Doing so satisfies as many people's *procedural claims* as possible, while maintaining the number of *benefit claims* that are

satisfied. The inverse is true when the number of claimants decreases, as happens when the benefits being distributed are urgent life saving treatments that are not provided before some claimants die. In cases where the number of claimants decreases over time and there are no other morally relevant changes, then one should run the lottery as soon as possible to satisfy as many people's *procedural claims* as possible.

When the quality of benefits that beneficiaries have a claim on increases over time (i.e. the benefits being distributed become better), all other things being equal, we ought to wait before distributing the benefits. For example, if we know that a more effective medical treatment will be legalised and available next week and there will be no disvalue in postponing treatment until the more effective treatment is available, then we ought to wait until we can administer the better treatment. That better satisfies their *benefit claims*[77]. Conversely, all other things being equal, if claimants' medical prognoses worsen if not treated immediately, as is the case with many medical conditions, we ought to administer treatments as soon as possible. Cases where only the number of claimants and available benefits change while outcome groups are of all equal size are relatively simple and the implications drawn from the dual structure view are fairly straightforward and intuitive. When the number of claimants or quality of benefits increases over time, we ought to postpone distributing the benefits. When the number of claimants or the quality of benefits decreases over time, we ought to distribute the benefits as soon as possible. The more

[77] Assuming they have a claim on the better benefit or the best possible benefit.

complex cases are when the requirements of *benefit claims* and *procedural claims* conflict, as in the real world case of corneal transplantations I will discuss next.

Consider the transplant of corneas to those who have suffered degeneration of their eyes. In the real world, both the number of claimants and corneas available for transplant increase over time, but there are many more people that could benefit from such a transplant than can be treated. Without treatment, the degeneration of their eyes continues and their prognoses can get worse. It is an important, but not urgent treatment. How should we act in those cases?

The procedurally fairest thing to do is to give as many claimants as possible a fair chance of benefiting. Doing so satisfies as many of the *procedural claims* as possible. In order to do so, one would wait to run the benefit distributing lottery until the maximum number of claimants was available at one time. This involves not treating people immediately, but waiting for additional people to develop optical degeneration before administering the lottery. However, it would not be fair to postpone the benefit distributing lottery indefinitely, because donor corneas are only transplantation viable for a limited period of time – that is, they expire. That the corneas expire is relevant for their procedural claims as no one is treated fairly when no one benefits and the benefits go to waste[78]. Instead, fairness requires us to postpone the benefit

[78] Here, I assume a non-comparative, absolute view of fairness. On such a view, fairness does not merely require comparatively *equal* satisfaction of equally strong claims, so that giving all claimants a 0% chance of benefiting is equally fair as giving all claimants a 100% chance of

distributing lottery until just before a cornea is no longer transplantation viable. However it is clear that is *not* the *right* thing to do, especially since the quality of the benefits worsen over time. Another benefit of the dual structure view is that it can be used to explain this. To see this, we will first have to consider how to satisfy the *benefit claims* in the above case.

The *benefit claims* are best satisfied by giving as many potential beneficiaries the best benefit that they have a claim on[79]. In the abovementioned corneal transplantation case, when there are many more claimants than benefits and the quality of benefits decrease over time, potential beneficiaries' benefit claims require us to distribute the benefits as soon as they become available. Just as maximising the satisfaction of *procedural claims* is not clearly right, it is not clearly right that we ought to maximise the satisfaction of *benefit claims* in these kinds of scarce distributional cases. The *dual structure view* explains this - we have two sets of claims and maximally satisfying only one set of claims is only

---

benefiting. Instead, fairness requires giving potential beneficiaries what they have a claim on and in cases of scarcity, best approximating the fairest situation where all claimants get a 100% chance of benefiting, where this is understood in absolute terms. For the purposes of this paper, I remain neutral on which is the best approximation (ie the most fair) of the fairest situation. I provide a more substantive account of a non-comparative view of fairness in my 'Weighing up Weighted Lotteries'. For a recent treatment of the comparative/absolute distinction, see Larry Temkin's 'Justice, Equality, Fairness, Desert, Rights, Free Will, Responsibility, and Luck' in Carl Knight and Zofia Stemplowska, *Responsibility and distributive justice,* Oxford University Press, 2011)

[79] One can imagine cases in which the benefits that people have claims on are accompanied by extra benefits that they do not have claims to. For example, imagine that the life saving buoys that are available are manufactured by different manufacturers. One manufacturer runs a special, unexpected promotional campaign that rewards people saved through the use of their life buoys in a limited time period with a special luxury holiday, which the manufacturer then uses in their advertisement campaigns. Assuming that such a promotional holiday is a benefit rather than a burden, it is a benefit they have no claim to and as such are not relevant to satisfying their *benefit claims*. As such, whether or not such a benefit gets better or worse is irrelevant to the satisfaction of their *benefit claims*, though it may be relevant for what we ought to do all-things-considered.

possible in the above case at the expense of the other. What we ought to do involves a trade off between the *procedural claims* and the *benefit claims*. Such a trade off, if it is to give any weight to both sets of claims (as is intuitively plausible in the corneal transplantation distribution case), will involve waiting some amount of time so that more claimants can be given a fair (positive) chance of benefiting despite some worsening prognoses. Precision regarding how much time depends on both axiological and empirical factors. The axiological factors include the value of the benefits distributed, the importance of fairness and the strength of beneficiaries' claims. The empirical factors include the speed and the amount the benefits of corneal transplantation reduce over time; when and how many claimants develop new *benefit claims* on corneal transplants; and when and how many new benefits (donor corneas) become available and *expire*. Settling all of these issues is beyond the scope of this paper, but given the large value of receiving a corneal transplant I expect that the right trade off involves some waiting, resulting in limited worsening of prognoses but no expiry of donor corneas.

What is not beyond this paper is a new proposal about how lotteries satisfy potential beneficiaries' claims on benefits. The dual structure view I have advocated posits the existence of two sets of claims in such distribution cases - *benefit claims* and *procedural claims*. In cases where the benefits are scarce and indivisible, appropriately run lotteries give potential beneficiaries a fair chance of benefiting, thereby satisfying their *procedural claims* while leaving their *benefit claims* unsatisfied until they actually receive the benefit. My dual

structure view parallels the structure evident in the *Chancy Doctor* case, while fulfilling all of the theoretical desiderata set out in Section I. It does so by giving theoretically satisfying explanations of temporally extended cases, such as *Extra Buoy* and *Stormy Seas*, as well as cases where only partially fair lotteries are run, while still avoiding *iterative absurdity*. The dual structure view is not only of theoretical interest, but can be used to shed light on practical problems of scarce benefit distribution. In particular, it has implications for temporally extended cases where the number or quality of available benefits or potential beneficiaries change over time, as occurs in a number of medical cases, such as the distribution of corneal transplants. As such, my proposal not only solves a problem of meta-ethical interest but also has important implications for the implementation of lottery distribution mechanisms in real world cases.

Bibliography

Broome, John. *Ethics Out of Economics.* Cambridge: Cambridge University Press, 1999.

Broome, John. "Uncertainty and Fairness." *The Economic Journal* 94, no. 375 (1984): 624-632.

Hirose, Iwao. "Weighted Lotteries in Life and Death Cases." *Ratio* 20, no. 1 (2007): 45-56.

Knight, Carl, and Zofia Stemplowska. *Responsibility and Distributive Justice.* Oxford University Press, 2011.

Kumar, Rahul. *Consensualism in Principle: On the Foundations of Non-Consequentialist Moral Reasoning.* Routledge, 2001.

Lang, Gerald. "Fairness in Life and Death Cases." *Erkenntnis* 62, no. 3 (2005): 321-351.

Otsuka, Michael. "The Fairness of Equal Chances." (Unpublished Manuscript).

Parfit, Derek. *On what Matters.* Oxford University Press, 2011.

Peterson, Martin. "The Mixed Solution to the Number Problem." *Journal of Moral Philosophy* 6, no. 2 (2009): 166-177.

Rakowski, Eric. "Taking and Saving Lives." *Columbia law review* 93, no. 5 (1993): 1063-1156.

Saunders, Ben. "A Defence of Weighted Lotteries in Life Saving Cases." *Ethical theory and moral practice* 12, no. 3 (2009): 279-290.

Timmermann, Jens. "The Individualist Lottery: How People Count, but Not their Numbers." *Analysis* 64, no. 2 (2004): 106-112.

Valuing the Few over the Many[80]

Consider the following Scanlonian *transmitter case*[81]:

> Jones has suffered a serious accident in the transmitter room of a television station where electrical equipment has fallen on his arm. While this accident is not fatal and will cause no permanent damage, until Jones is saved he will suffer from serious pain due to electrocution and lack of treatment. It is not possible to rescue him without turning off electricity to the transmitter room for fifteen minutes, thereby interrupting television transmission for that length of time. Amongst other programs, a live World Cup match is in progress, watched by many people and the game will continue for over an hour (the 2010 FIFA World Cup had a total cumulative television audience of over 3.2 billion viewers, 46.4% of the global population)[82].

In this case, as Scanlon notes, our moral intuition is that we should not wait to save Jones no matter how many viewers there are. This is the true even if we note that some fanatic viewers enjoy watching the televised World Cup game very much. In addition to this deontic moral judgment, some have the

[80] I am grateful to comments from John Broome, Krister Bykvist, Daniel Halliday and Derek Parfit on earlier versions of this paper.
[81] Thomas Scanlon, *What we owe to each other,* (Cambridge, Mass. ; London: Belknap Press of Harvard University Press, 1998), 235-36.
[82] "FIFA Media Release: Almost half the world tuned in at home to watch 2010 FIFA World Cup South Africa," [cited 2012]. Available from http://www.fifa.com/worldcup/archive/southafrica2010/organisation/media/newsid=1473143/index.html.

axiological judgment that it is better to save Jones than to wait, no matter how many viewers there are. As I will show, this axiological judgment is inconsistent with some widely held, prima facie plausible claims about value. This paper will evaluate different attempts to incorporate these axiological judgments into moral theory despite the aforementioned inconsistency and defend a particular view about how to vindicate these judgments. To this end, the paper is structured as follows. The first section responds to a number of objections to the features of Scanlon's original case and outlines the type of distributive case this paper will focus on. I will argue that our axiological judgments in these cases are robust over both interpersonal and intrapersonal cases. The second section will show how our axiological judgments in these cases are inconsistent with some widely held claims about value and will proceed to evaluate a number of suggestions about how to incorporate these axiological judgments into a plausible theory of value. The final section will respond to some objections to our intuitive axiological judgment.

Section I - Disparate Conflict Cases & The Second Standard View

The *transmitter case* is a particular case of a general distribution problem that I will focus on in this paper, namely cases when one can either benefit a small, sometimes singleton, group of potential beneficiaries a large amount or benefit a much larger number of potential beneficiaries a very small amount. Call these cases *disparate conflict cases*. I will call the small group of potential beneficiaries

'*the few*' and the large group of potential beneficiaries '*the many*'. Many have the judgment that in cases where the disparity between the quality of the benefits given to the *few* compared to the *many* is sufficiently large, that it is better to benefit the *few* rather than the *many* no matter how numerous the *many* are. This judgment is robust – I will argue that it holds over a variety of cases involving both benefits and burdens.

Firstly, I will consider some objections to particular features of Scanlon's version of the case and show that our intuitive axiological judgment does not depend on those features. Kamm has objected to the *transmitter case* on the basis that it implicitly invokes the moral importance of causing harm[83]. Jones' suffering is caused by the electrical equipment used in the television broadcast that is benefiting many people. If we choose to continue the television transmission, it is plausible to characterize the case as one in which we are causing the suffering of Jones in order to benefit many people. Many moral philosophers believe that, all other things being equal, not-harming takes priority over benefiting and this could be used to explain our intuition in the above case. Furthermore, as Parfit notes[84], the example stipulates that Jones is in a much worse starting position (extreme pain) than the television watchers, which may also affect our intuitions in these cases. In order to show that these factors are not necessary to produce our intuition that it is better to benefit *the few* over *the many*, consider the following cases.

[83] Frances Myra Kamm, *Intricate ethics : rights, responsibilities and permissible harm,* (New York, N.Y. ; Oxford: Oxford University Press, 2007), 471-473.
[84] Derek Parfit, *On What Matters,* (Oxford: Oxford University Press, 2011), Volume 2, 193-212.

Imagine we can provide a benefit B_1 to one person (*the few*) or some qualitatively lesser benefit B_2 to *m* people (*the many*). Stipulate that prior to receiving the benefit all people are equally well off - of the same age, have equal life expectancy and have had exactly the same quality of life and so on. Suppose the benefit of B_1 is extending the beneficiary's life by ten years, during which they live a happy life. Suppose the benefit of B_2 is extending the beneficiaries' life by ten seconds, during which they live an equally[85] happy life. Here, many judge that it is better to provide B_1 to a single individual instead of giving B_2 to any number of individuals, even when the total length of life extended through B_2 benefits is greater than ten years. This intuition can be reproduced with other benefits that are qualitatively disparate – imagine B_1 is the benefit of experiencing bliss for a year and B_2 is the benefit of experiencing extremely minor pleasure (e.g. finding a drink you want to purchase is slightly cheaper than you expected) or B_1 is the benefit of preventing the beneficiary from suffering from a very significant harm (e.g. permanent quadriplegia) and B_2 is the benefit of preventing the beneficiary from suffering a very minor harm (e.g. temporary sore throat or hangnail). Again, stipulate that in these cases that there is no other morally relevant difference between potential beneficiaries. Not only do we have this axiological intuition about what is the better outcome, we also have a deontic intuition about which outcome we ought to pursue. Intuitively, we ought to benefit *the few* and not the *many.*

[85] That is, their quality of life during the 10 seconds is equal to the quality of life of 10 seconds of the individual benefited by B_1.

Furthermore, these two intuitions (deontic and axiological) are not restricted to benefits but also apply analogously to harms. Imagine you have only two morally bad choices: causing extreme harm to a single innocent person (e.g. quadriplegia or severe psychosis) or a minor harm to many innocent people (e.g. a mild minute long headache). Again we have the intuition that it is worse to subject the individual to the major harm than any number of people to the minor harm and given we only have two choices, that we ought to harm the many a small amount instead of harming the individual a great deal.

Not only do these intuitions apply to interpersonal cases, they also apply to single person cases. It is intuitively worse to cause an extreme harm to a single person, such as month of torture, than to inflict a qualitatively minor harm like a mild headache to them, no matter how long that headache lasts. Similarly, it is better to give someone an extremely large benefit (e.g a month of deserved bliss or prevention of a severe burden such as quadriplegia) than a qualitatively small benefit like a minor pleasure, no matter how long that small benefit lasts. Intuitions of this kind are held by both consequentialists [86] and nonconsequentialists[87] alike.

[86] Carlson affirms both the deontic and axiological versions of the intuition, whereas Mayerfeld affirms the deontic version but is not clear about the axiological version. Erik Carlson, "Aggregating Harms - Should We Kill to Avoid Headaches?" *Theoria* 66, no. 3 (2000): 246-255.Jamie Mayerfeld, *Suffering and Moral Responsibility,* (New York ; Oxford: Oxford University Press, 1999), 182-187.

[87] Scanlon affirms the deontic version of the intuition and a weaker version of the axiological intuition. Scanlon claims that always acting to benefit *the few* in the actual world would not lead to consequences sufficient to reasonably reject the principle to always benefit *the few* as opposed to *the many*, but not the stronger claim that I do, namely that we have the axiological intuition that the outcome is better when we benefit *the few* rather than *the many* in disparate conflict cases no matter how outnumbered *the few* are. Scanlon, *What we owe to each other,* 234-238.

Our axiological interpersonal intuitions in cases like this are largely captured by Temkin's *Second Standard View* (*SSV*) —"*Trade-offs between Quality and Number Are Sometimes Undesirable Even When Vast Numbers Are at Stake*: If the quality of one kind of benefit is "sufficiently" low, and the quality of another kind of benefit is "sufficiently" high, then an outcome in which a relatively small number of people received the higher quality benefit would be better than one in which virtually any number of (otherwise) similarly situated people received the lower quality benefit."[88] Temkin also believes that the view also applies to harms, so when some harms are sufficiently qualitatively severe, it is worse that they occur than virtually any number of sufficiently qualitatively minor harms. For brevity, I will refer to sufficiently qualitatively low benefits (or harms, respectively) to which the Second Standard View applies as *minor benefits/harms* and sufficiently qualitatively high benefits (or harms, respectively) to which the Second Standard View applies as *major harms*.

I claimed that Temkin's view *largely* captures the preceding intuitive judgments. It does not precisely capture them, as Temkin's formulation is weaker than the intuitions. This is because he adds the qualifier 'virtually' to 'any number', so he leaves it open that there is some very, very large number of minor benefits that would be better than a small number of major benefits. In a footnote[89], he makes it explicit that he is sympathetic to the position that drops the 'virtually' qualifier. He claims that he "cautiously" adds the qualifier that makes the view

[88] Larry Temkin, *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*, Oxford University Press, 2012), 32.
[89] ibid., 33n14.

weaker so that his argument "will have force for as many people as possible". Given Temkin's sympathies, and the fact that inclusion of this qualifier does not capture the intuitive judgments previously mentioned in this section, for the purposes of this paper, I will amend Temkin's stated Second Standard View. When I refer to the Second Standard View in this paper, I will be considering it without the 'virtually' qualifier.

The paper will not offer an account of precisely which cases the Second Standard View applies to, nor is one necessary for the purposes of the paper. For the purposes of this paper, all I require is that there are some benefit distribution cases to which the Second Standard View applies. In this paper, I want to argue that the view can be plausibly incorporated into value theory despite its inconsistency with some widely held views about how to aggregate the value of benefits and harms. My suggestion about how to resolve this inconsistency is independent of which views one takes on which cases one takes the Second Standard View to apply to, as it can be adopted on a number of different views about which benefits are sufficiently low or sufficiently high.

However before illustrating the axiological inconsistency and evaluating a number of proposed solutions to it, I will defend the Second Standard View from some prominent objections in the next section.

Section II - Defending the Second Standard View

Even though the intuitions that support the Second Standard View are widely held and robust over benefits and harms, they are subject to two kinds of objections. I will outline and respond to those objections in this section. One set of our objections claim that the intuitions in support of the Second Standard View are inconsistent with more strongly held intuitions and as such we should reject the Second Standard View. The second type of objection attempts to show that the intuitions supporting the Second Standard View can be shown to be independently unreliable. I will defend the Second Standard View and the intuitions that underlie it from these two objections in turn. It would not be worth seeing how to best incorporate the Second Standard View into an axiological theory if there are decisive arguments showing that the intuitions underlying it are unreliable or objectionable.

Norcross claims that one particular intuition underlying the Second Standard View conflicts with more strongly held intuitions and therefore should be rejected[90]. The intuitive judgment he targets, which he refers to as *Worse*, is that all other things being equal, it is worse that one person suffer a premature death than any number of people suffer moderate temporary headaches for twenty four hours. This conflicts with three other intuitively supported claims, which he calls *Risk*, *Many Headaches* and the transitivity of the 'better than'

[90] Alastair Norcross, "Great harms from small benefits grow: how death can be outweighed by headaches," *Analysis* 58, no. 2 (1998): 152-158.

relation. *Risk* is the claim that other things being equal, it is better that one person die than five million each incur a one in a million risk of dying. He supports this axiological claim with a deontic intuition. He asks reader to consider what is intuitively right in a case where we can either send a poisonous gas into room R_1 or room R_2, where sending it into R_1 will cause one person to die and sending it into R_2 will put five million people at a one in a million risk of dying. I will not dispute *Risk*. *Many Headaches* is the claim that other things being equal, it is better that five million people each incur a one in a million risk of dying than they each suffer a moderate headache for twenty four hours. He supports this axiological claim by first considering a single person who subjects himself to a one in a million chance of death by driving to the pharmacy to purchase an analgesic that will prevent him from suffering a moderate headache for twenty four hours. That act is rational and permissible. He assumes that if that analgesic seeking act is rational and permissible, it is better to subject one to that risk of death than have the moderate headache. I will deny this assumption. In order to support *Many Headaches*, he then asks us to imagine five million people each undergoing this pharmacy trip (each in a different city to prevent complications) and finds it to be intuitively rational. As Norcross notes, *Risk*, *Many Headaches* and the transitivity of the 'better than' relation entail the denial of *Worse*.

While there are reasons to reject the claim that 'better than' is transitive, Norcross's argument can be blocked independently of this intransitivity claim. This is because Norcross's analgesic case does not support *Many Headaches*

because of a morally important distinction between his case and the *Many Headaches* claim. It is morally relevant that the potential headache sufferer consents to the risk she puts herself under when she drives to the pharmacy – this is what makes the action permissible[91]. I will argue that when there is no consent, it is both impermissible and worse to subject someone to a one in a million chance of death instead of a moderate 24 hour headache. To see this, consider the following case, when you have to choose to press one of two differently coloured buttons. The red button causes one person to have a moderate headache, whereas the blue button causes the same person to incur a one in a million risk of death (though if they do not die, they do not suffer any adverse affects). If we do not have the person's consent to subject them to risk of death, my intuition is that the right thing to do is choose the red button – it would be intuitively wrong to press the blue button. Furthermore, it is intuitively worse for the person to subject the person to risk of death without their consent than to give them a headache. However, if we have the person's consent to subject them to risk of death in order to avoid the headache, then my intuition is that it is permissible to choose the blue button.

The relevance of consent for permissibility can be further brought about if we consider that the recipient of the harm is the button pusher, so they have to choose between two self-inflicted harms. In such a case it is intuitively both

[91] It would also be intuitively rational and permissible for the person to decide to stay at home and suffer the headache and not risk their own life.

permissible and rational for them to pick either option[92]. Against Norcross I argue that the mere fact that the headache relief option is rational and permissible does not entail that it produces a better outcome than its alternative. After all, both actions are intuitively rational and permissible. Endorsing the view that the rationality and permissibility of an act entails that act is better than its alternative(s), and that both acts are rational and permissible, leads to a contradiction - that headache relief is better than a one in a million chance of death and one in a million chance of death is better than headache relief. We ought to give up the view that the rationality and permissibility of an act entails that it is better than its alternative(s). It is both rational and permissible for me to drink still water now that I am thirsty, but that does not entail it is better than some other drink.

Norcross's example involves self consent, which makes it more analogous to the self-harming button pusher case than conflict cases when you are deciding what to do to others. *Disparate conflict cases* by definition always involve other people because they always contain large numbers of claimants and as such the individual choosing who to benefit cannot provide such consent for all claimants. As I have argued, the axiological claim, *Many Headaches*, does not follow from his claims about the rationality and permissibility of his analgesic seeking drivers. Furthermore, in *disparate conflict cases* where we do not have the consent of all

[92] Here I do not make the stronger and more contentious claim that self consent makes any self inflicted harm permissible. Many, if not all, libertarians would be concerned to defend this stronger claim as long as the self inflicted harm did not harm any other individuals. Here, I only defend the claim that such consent makes minor harms (such as very small risks of premature death) permissible that would be impermissible without such consent.

potential beneficiaries, my intuitive judgment is that it is *worse* to subject five million to a one in a million risk of death than to subject those same five million to temporary moderate headaches. Norcross fails to provide a sound argument for *Many Headaches*, and we have intuitive reason to reject it, so we ought to reject it and with it, conclude that his argument against *Worse* is not sound.

Furthermore, it should also be noted that Norcross's failed argument is only an argument against a particular instance of the Second Standard View. In order to effectively show there were a set of intuitions that were more strongly held than those that underlie the Second Standard View, he would have to construct analogous cases for other intuitions. Thus even if Norcross's argument was sound, it would not undermine the Second Standard View's application to other cases. For the purposes of this paper, I only need the Second Standard View to apply to some distributive cases. I believe it applies to *Worse*, but even if it does not, there are other cases such as those described in Section 1 with which to motivate the Second Standard View. I now turn to another attempt to show that the intuitions that support the Second Standard View are inconsistent with other more strongly held intuitions.

Broome appeals to the UK National Health Service's (NHS) actual practices to argue against irrelevant utilities. While Broome wrote this paper prior to Temkin's coining of the Second Standard View, his argument could be construed

as an objection to that view[93]. When patients in NHS hospitals suffer from minor pains they are given analgesics. The cost of these analgesics adds up over the years, in such a way that eventually they cost as much as it would to save someone's life. Broome thinks that the NHS's distribution of benefits is intuitively right. The NHS's distributive procedures could be viewed as being inconsistent with the *Second Standard View* because it suggests that the prevention of some number of qualitatively minor harms outweigh the value of saving someone's life.

My response is two fold. Firstly, current government policy at best provides defeasible evidence in support of a moral claim. In many historical cases, government policy was immoral - consider slavery, the denial of women's suffrage and indigenous populace oppression. We should therefore be wary of the morality of government policy. This is especially the case when the moral claims in question are only of recent origin. It is to be expected that claims advocated in moral and political philosophy take some time to be implemented. Many of the liberties that Mill defended in *On Liberty* in 1859[94], such as freedom of speech, were only implemented effectively (and even then, only in some developed countries) in the 20th century. The *transmitter case* and its implications are products of the 1990s and still controversial and as such it would be unreasonable to expect these developments to be reflected in public

[93] Broome, J. "All Goods are Relevant" *Summary Measures of Population Health: Concepts, Ethics, Measurement and Applications,* eds. Christopher Murray, Joshua Salomon, Colin Mathers and Alan Lopez (Geneva, Switzerland: World Health Organization, 2002), 728. Broome wrote this paper prior to Temkin's coining of the *Second Standard View*.
[94] John Stuart Mill, *On Liberty* (Boston: Ticknor and Fields, 1863).

policy. As such, current government policy provides weak, defeasible evidence against the Second Standard View, evidence which can be outweighed by the defence of the intuitions I provide in this paper.

Secondly, I do not share Broome's intuitions that what the NHS is doing is right, given the lack of medical resources and inability of the NHS to prevent all preventable cases of premature death (and other major harms) due to lack of medical resources. I do not deny that the NHS should make analgesics available – my intuitive judgment is that the NHS should sell these analgesics to patients, perhaps at cost price, in order to direct such limited funds to treating more severe harms such as premature death. This seems to me better than allowing people to suffer severe harms that could have otherwise been prevented. I have responded to objections that claim that the intuitions underlying the *Second Standard View* conflict with more strongly held intuitions and now I turn to Broome's objection that intuitions underlying the Second Standard View are unreliable.

Broome objects to trusting our intuitions in cases that involve large numbers[95]. Broome notes some humans cannot grasp large numbers and as such our intuitions involving such numbers are unreliable. This argument only draws doubt on the reliability of intuitions in such cases; it does not provide decisive evidence against them. Broome supports his claim with an example - some people have the wrong intuition about evolution, namely they have the intuition

[95] John Broome, *Weighing lives*, (Oxford: Oxford University Press, 2004), 57-58.

that the process of natural selection, no matter how long that process goes on for, could not lead from primordial slime to intelligent life. However, like there are people who can grasp the truth of evolution intuitively, there are those who can grasp large numbers, including infinities, such as theoretical mathematicians. If it was the inability of people to grasp large numbers that caused ordinary people to have intuitions in support of the Second Standard View, we would expect that theoretical mathematicians would not share our axiological intuitions in these large number cases. However, all of the mathematicians I have asked about disparate conflict cases agree that there are some benefits that are disparate enough that it is better to benefit *the few* rather than *the many* no matter how *many* there are. Given this evidence, it is not clear that it is our inability to grasp numbers that explains our axiological judgments in disparate conflict cases.

Having responded to objections to the deontic and axiological judgments underlying the Second Standard View, I will now proceed to show the Second Standard View's inconsistency with a plausible and widely held view about how to aggregate the value of benefits and harms.

Section III - Resolving Axiological Inconsistency

This Second Standard View conflicts with widely held and intuitively supported axiological aggregation principles, namely the following set of claims[96]:

[96] Both of these claims have a great deal of intuitive plausibility and have been assumed by many philosophers for hundreds of years. For example, Bentham writes "The community is a fictitious

(1) In calculating the disvalue of a state of affairs, the disvalue of different individuals' harms should be aggregated in a particular way, namely that the moral disvalue contributed by one or more harms of qualitative level *x* is directly proportional to the number of people suffering from that harm of *x* qualitative severity. This form of aggregation ensures that when harms are equally qualitatively intense, each person's harm is of equal consideration in calculating the disvalue of a state of affairs.

(2) All harms, whether severe (e.g. agonizing torture) or minor (e.g. hangnail) are of finite disvalue in all contexts.

If these prima facie plausible claims are all true, one can produce a conclusion that conflicts with the Second Standard View as some (perhaps extremely large) number of people suffering from some extremely minor harm is worse than (that is, can outweigh the disvalue of) one person suffering from a very severe harm. This is because the value of preventing some finite number of minor harms can be aggregated in a way such that there is some number of minor harms prevented that is better than preventing a single severe harm. In fact, together these claims entail that for any finite harm *x* and any number *n*, and some qualitatively lesser harm *y*, there is some number *m* such that it is worse

---

body, composed of the individual persons who are considered as constituting as it were its members. The interest of the community then is, what?—the sum of the interests of the several members who compose it... It is in vain to talk of the interest of the community, without understanding what is the interest of the individual. A thing is said to promote the interest, or to be *for* the interest, of an individual, when it tends to add to the sum total of pleasures." Jeremy Bentham, *An Introduction to the Principles of Morals and Legislation (Collected Works of Jeremy Bentham)*, Clarendon Press, 1996), Chapter 1, iv.. I will not argue for either (1) or (2), but will assess attempts to reject either (1) or (2) in order to vindicate the Second Standard View.

that *m* people suffer *y* than *n* people suffer *x*. To maintain the Second Standard View, (at least) one of these two claims must be rejected. I will consider the plausibility of rejecting each of these claims in order to evaluate the different ways in which the Second Standard View can be theoretically vindicated and I will ultimately endorse a particular rejection of (2).

Firstly, we might deny (1), as Taurek does[97]. Taurek thinks that in a case where all relevant people are morally alike in all relevant respects, five unrelated people suffering a premature death is not morally worse than one person suffering a premature death because both are equally bad when considering the harms to any individual. Five unrelated individuals each dying does not add up or entail any individual experiencing a loss greater that is in any way greater than the loss suffered by any one of the five. No one can claim that it is worse for him or her to die than it is for others[98]. Taurek's conclusion is stronger than the denial of (1), as it not only denies the particular kind of aggregation advocated in (1), it denies any kind of axiological aggregation in interpersonal conflict cases where no one is related. That is, it is not worse, according to Taurek, that more unrelated people die prematurely than fewer people die prematurely.

[97] John Taurek, "Should the Numbers Count?" *Philosophy and Public Affairs* 6, no. 4 (1977): 293-316.
[98] Taurek does recognize the moral importance of greater losses in other cases – if we had the choice of protecting five objects, each of value *z* to us, and one object of value *z* to us, we would be justified in protecting the larger number because the larger group of objects is five times as important to us than one object.

Partially because of its startlingly implausible conclusion, Taurek's claim is much discussed. Rather than evaluate all objections to Taurek's view, I will give what I take to be compelling objections to Taurek's conclusion.

While I accept Taurek's ontological claim that there is no individual that experiences a loss that is greater than the loss suffered by any one of the five, I deny that such an entity would have to exist in order to justify the claim that it is worse for five to die prematurely than one die prematurely. I will argue that there is no plausible general ontological requirement for the relata of evaluative relations nor a specific ontological requirement that applies in benefit distribution cases like this one.

In general, we do not require the relata of better than or worse than relations to exist in order to evaluatively compare them. We can and regularly do evaluate merely possible options that never come into existence. To take a common example, I can evaluatively rank main courses on a restaurant menu in order to select which dish I want to order, even though only one particular dish comes into existence after I order it (perhaps two, if I'm greedy!)[99]. Consider another example. One of the central explanations of why premature death is bad for the one that dies is the loss of the valuable future that they would have had if they had not died. However in the case where the person dies, those goods never actually exist and are just possibilities. It would be implausible to require that all relata of evaluative relations exist.

[99] The mere existence of the recipe or the ingredients for the dish, do not entail that the actual dish ordered exists.

Not only is there no general principle that requires the relata of evaluative relations to exist in order to make true claims about whether the relation holds, I see no cogent argument from Taurek that explains why the particular individual that Taurek claims does not exist is required to explain why five dying is worse than one. In order to explain why it is worse for five to die than one to die, one need not appeal to the existence of a special high level entity that experiences the badness of all five premature deaths. Instead one can appeal to the plurality of evaluative facts about individuals and make facts about the relative badness of outcomes reducible to those individual facts[100].

However for the purposes of this paper, there is a more important reason to reject Taurek's view as a way of vindicating the intuitions enshrined in the Second Standard View. The claim that it is not worse for more unrelated people to suffer a harm rather than less is startlingly implausible - much more implausible than the denial of the Second Standard View. The striking implausibility can be brought out by changing the numbers of the case. Taurek's view implies denying that it is worse that one million unrelated people suffer from agonizing pain than one person suffering from that same agonizing pain. This is an extremely implausible view. Our belief that it is worse for the million is much stronger than the strong support for the Second Standard View. If vindicating the Second Standard View required denying this claim, it would give us compelling reason to deny the Second Standard View. Fortunately, there are

[100] Ralf Bader defends a view of this kind in Ralf Bader, "Neutrality and Conditional Goodness," (Unpublished Manuscript).

more plausible ways of vindicating the Second Standard View. To these other ways I now turn.

Mayerfeld has noted that it is somewhat less counter-intuitive to deny (1) only for qualitatively minor harms such as those suffered by *the many*[101]. That is, he suggests that qualitatively minor harms cannot be summed in a way such that more numerous minor harms of qualitative severity *y* are worse than less numerous minor harms of qualitative severity *y*. This is a weaker claim than Taurek's. As we saw, Taurek's view denies (1) for *all* harms, no matter how severe, as long as those harms happen to morally relevantly similar individuals.

Despite its relative weakness, Mayerfeld's suggestion is still rejectably counter-intuitive. Fifty people suffering from a mild headache is intuitively worse than a single person suffering from a mild headache of the same intensity. Denying this is much more intuitively costly than denying the Second Standard View. Furthermore, it is difficult, if not impossible, to locate the qualitative level that distinguishes harms that should be summed according to (1) from those which should not. There appears to be no non-arbitrary qualitative level that could plausibly serve this role. This is not merely a practical problem that those intending to implement Mayerfeld's suggestion would encounter. The fact there appears no non-arbitrary level with which to implement Mayerfeld's denial of (1) indicates that it is a metaphysical and not merely an epistemological or

[101] Mayerfeld, *Suffering and Moral Responsibility*, 187.

practical problem. As a result, both Taurek's and Mayerfeld's suggestions about how (1) may be rejected should themselves be rejected.

Affirming the lexical priority of qualitative considerations over quantitative considerations is another way to deny (1). This is evident in leximin, which states that it is always lexically better to benefit the worst off (in this case, the most severely harmed) over any number of others who are better off[102]. Leximin, however, has the deeply troubling implication that for any two qualitatively different harms, all other facts about the relative welfare of potential beneficiaries equal, it is always better to benefit the more severely harmed than the less severely harmed. The trouble with this is most evident in harms that are qualitatively similar in the severity of burden to each individual, but the numbers of people suffering the less severe harm far outnumber those suffering the more severe harm. For example, the view claims that it is better to cure one person of quadriplegia instead of a million (or any number of!) people of paraplegia. This implication is highly implausible and as a result leximin should be rejected.

Mayerfeld suggests we can deny (1) by arguing that at some point the addition of another harm suffered by an additional individual ceases to contribute to the moral disvalue of a state of affairs[103]. To illustrate by example, the moral

[102] While leximin denies summation, it gives some consideration to quantitative considerations because if one could prevent 100 people from premature death or 20 different people from premature death, all other things being equal, it would claim it is better to save the greater number as the qualitative harms are the same. ibid., 186-187.
[103] ibid., 187.

disvalue of people suffering from some minor harm, such as a hangnail, is worse for each individual suffering that hangnail until some number of people harmed is reached (call this number *the non-contribution point*), so that any additional individuals do not affect the moral disvalue of the state of affairs. This number can be selected so that the moral disvalue of a severe harm can never be outweighed by any number of minor harms. This is because after the *non-contribution point* further harms of that same level do not affect the state of affairs' moral disvalue.

I can see two possible variants of this view - one view would claim that *non-contribution points* apply to all qualitative levels of harm, one view would claim that they only apply to some types of harms. The former view is more parsimonious. Furthermore, in order to vindicate the Second Standard View, the latter view would have to ensure that the minor harms that the *few* suffer are included in the types of harms that have *non-contribution points* and that the distinction between the two different types of harms reflected some genuine moral difference. I will not consider each of these two views separately, as they are both subject to the same decisive objections.

Mayerfeld's denial of (1) has deeply counter-intuitive implications and has serious difficulties motivating precise *non-contribution points*. Two significant counter-intuitive implications immediately come to my mind. Firstly, it violates our intuition that each moral agent should be of equal consideration in the moral evaluation of a state of affairs. According to the *non-contribution point*,

all harms that happen to people beyond the non-contribution make no difference to the disvalue of that state of affairs. Secondly, it claims that in cases where *the non-contribution point* is reached for some particular level of qualitative harm, the moral disvalue of the state of affairs is insensitive to the addition of more people being harmed at that same qualitative level. That is, assuming a hangnail's *non-contribution point* is 10,000, Mayerfeld's suggestion makes the implausible claim that it is not worse that 100,000 people suffer hangnails than 10,001 people suffer hangnails[104]. Furthermore, it is difficult, if not impossible, to determine where *the non-contribution point* is for qualitatively different harms. If one attempted to rely upon our intuitions to determine where *the non-contribution point* is located for any harm, one would face the second of the two problematic counter-intuitive results mentioned above – for any qualitative level of harm, it would always seem that there is some greater number of people suffering the same qualitative level of harm that is worse than some smaller number of people suffering that same qualitative level of harm. This suggests that there is *never* a *non-contribution point*. For these reasons, we should reject *the non-contribution point* attempt to theoretically vindicate the Second Standard View.

Carlson[105] suggests an initially more plausible denial of (1) by arguing that the marginal disvalue of additional harms of the *same qualitative severity* decreases

[104] This is particularly problematic for views that claim all qualitative levels of harm have non-contribution points. Such views would entail that there are non-contribution points for qualitatively serious harms like premature death and agonizing torture. Such views have counter intuitive implications almost as dire as Taurek's view. For such views, once the non-contribution point for premature deaths has been reached, it does not axiologically matter that thousands more die.

[105] Carlson, *Aggregating Harms - Should We Kill to Avoid Headaches?*, 246-255.

depending on how many others are suffering from the same harm. That is, while it is morally worse that three people suffer from a headache of intensity *h* than one person suffering from a qualitatively similar headache, it is not three times worse and furthermore, the contribution of moral disvalue of the second person to the state of affairs is greater than the third person but less than the first. A mathematical example Carlson's theory is as follows, where *n* is the number of qualitatively equally severe harms suffered, where $x_1,\ldots,x_n$ represents these harms (ordered axiologically, not qualitatively, from worst to least bad), $u_1,\ldots,u_n$ represents the values of $x_1,\ldots,x_n$ and k is a constant in such a way that $0<k<1$.

$$\text{The value of } x_1,\ldots,x_n = \sum_{i=1}^{n} u_i \times k^{i-1}.$$

This marginal disvalue can be adjusted so that more qualitatively severe harms are given greater weight than qualitatively minor harms and then the value of a state of affairs can be calculated by summing these weighted values. To illustrate how such a theory can account for the Second Standard View while rejecting (1), stipulate that the value of a premature death of a twenty year old innocent person is -20,000 while five minutes of a mild headache is -1. The total disvalue of any number of such mild headaches cannot be lower than the axiological limit -1 x (1/(1-k)) determined by the aforementioned equation and assuming that k = 0.999, this limit equals -1,000. Therefore it is never as morally bad as one instance of the aforementioned premature death (-20,000). This does not mean that major harms cannot be outweighed by slightly less major

harms. To see this, suppose permanent paraplegia of a single innocent person is -2,000. Still assuming k = 0.999, the maximum disvalue of innocent people suffering permanent paraplegia is -2,000,000.

While this suggestion seems the most plausible of the responses considered so far, I have three objections against Carlson's suggestion. Firstly, because Carlson's theory only decreases the marginal disvalue of additional harms when there is more than one person suffering from a harm of a particular qualitative level, given a sufficient range of qualitative minor harms, it would be consistent with Carlson's theory that some (perhaps extremely large) number of different types of minor harms (e.g. minor sore throat[106], short headache, hangnail) is worse than a single qualitatively serious harm (e.g. premature death, paraplegia). Thus in its current formulation, Carlson's suggestion cannot account for such a qualitatively varied version of the Second Standard View[107]. This objection could be overcome by creating a more complex weighting calculus that decreased the weight of the marginal disvalue of qualitatively different harms in some relevant qualitative range of harms. However, responding by making the moral evaluation procedure even more complex leads to further practical and theoretical problems. The more complex the evaluation procedure, the more difficult such evaluations are to conduct. Furthermore, the more complex the calculation, the more theoretically

[106] Note that there can be a many different individuals each suffering a slightly qualitatively different harm, e.g. different lengths and different intensities of suffering.

[107] Carlson recognizes this problem. He writes: "If a sufficiently great number of harms of a certain type have already occurred, the contributive value of any future harms of this type will be close to zero. Hence, there is a risk that such future harms will be given too little weight, compared to harms of other types (or to different benefits)." ibid.: 249.

unsatisfying the theory would be – such an accommodation of our intuitions would likely be a gerrymandered contortion of the theory to fit the intuitions as opposed to being truly explanatory and reflective of morality[108].

Secondly, like any rejection of (1), Carlson's suggestion denies our strongly held intuition that all moral agents are of equal value and should therefore contribute equally to our moral evaluation of a state of affairs. It seems objectionable that the 101st person suffering a harm is less axiologically weighty than the 1st person suffering a harm of the same qualitative level. This is especially problematic when there are particularly many people suffering a harm (i.e. when approaching the limit) as any additional individual suffering a harm is of negligible moral significance. This leads to my third objection.

Carlson's theory has the unacceptable implication that when there are very many instances of a severe harm, an additional harm of the same high severity is less bad than the first instance of a minor harm. To see this, consider two possible worlds, one containing 50 billion and one people suffering from serious and agonizing mutilation and one containing 50 billion people suffering from serious and agonizing mutilation and one person suffering from a moderate headache for a day. According to Carlson's theory the marginal disvalue of the 50 billionth and first person suffering from agonizing mutilation is weighted so that it is very small, so that mutilation's disvalue is less than the disvalue of the very first unweighted headache harm. As such, Carlson's theory highly

[108] Michael Otsuka, "Saving Lives, Moral Theory, and the Claims of Individuals," *Philosophy & Public Affairs* 34, no. 2 (2006): 128.

implausibly suggests that the former of the two worlds is better than the latter. This is only one particular example of my objection as Carlson claims his theory applies to all harms from worst to least bad. For Carlson, if there are sufficiently many instances of severe harms like quadriplegia, premature death by decapitation and permanent blindness, the disvalue of additional people[109] suffering from these severe harms is less than the disvalue of the first person suffering a minor harm. I take my argument here to be decisive against Carlson's theory. I have evaluated the prominent attempts to vindicate the Second Standard View by denying (1) and now turn to evaluating ways in which we can reject (2).

One could reject (2) by arguing that whatever severe harm being suffered by the *few* in cases subject to the Second Standard View is of infinite moral disvalue. In order to vindicate the variety of harms involved in the Second Standard View, a vast variety of severe harms (e.g. premature death, quadriplegia, agonizing torture) would all have to be of infinite moral disvalue. Like Taurek's suggestion, this suggestion seems strikingly implausible. Why would it be *infinitely* bad for a single person to suffer agonizing torture? In addition to this striking counter-intuitiveness, the proponent is forced into a dilemma.

Either all harms are of infinite disvalue or not all harms are of infinite disvalue. I will first consider the claim that all harms are of infinite disvalue. One could argue that if some severe harms are of infinite disvalue, all harms are of infinite

[109] In fact, with sufficient numbers of severe harms, the marginal disvalue of the last few hundred severe harms may be outweighed by a single less severe harm.

disvalue. For any harm of infinite disvalue, it seems plausible that there would be a harm that is slightly less qualitatively severe so that there would be some number of slightly less qualitatively severe harms that would be worse than one of the slightly more severe harm. To take a particular example, it would be worse for five to suffer agonizing torture for an hour than one to suffer from slightly more agonizing torture for seventy minutes. This is analogous to the objection against leximin. In order to vindicate this claim, if the qualitatively more severe harm is of infinite value, I take it that the qualitatively less severe harm would also have to be of infinite disvalue. However, there could be some number of even qualitatively less severe harms that would outweigh that infinitely disvaluable harm and would therefore also have to be of infinite disvalue. Assuming that it is true of all harms that there is some number of slightly less severe harms that are worse than a smaller number of the more severe harms, the problem would iterate and so arguably all harms would be of infinite disvalue. This is a particularly counter-intuitive version of the view as not only *some* harms are infinitely disvaluable, but *all* are. Furthermore, if all harms are of infinite disvalue, including the minor harms suffered by the *few*, this view would no longer vindicate the *Second Standard View* as some number of very minor harms could evaluatively outweigh very major harms.

On the other horn of the dilemma, if not all harms are of infinite disvalue, then the theory will have an unacceptable implication. That is, it will deny that supposedly infinitely bad harms can be outweighed by some number of non-infinitely bad moderate harms. Take the qualitatively least bad harm of infinite

disvalue and then consider a harm slightly qualitatively less bad than that harm so that it is of finite disvalue. It seems intuitive that there is some number of the most severe finitely disvaluable harm that would be worse than one of the least severe infinitely disvaluable harm. This need not even be the most severe finite harm. To show how this horn of the dilemma is counter-intuitive, we judge that the moderate harm of someone breaking their arm is of significant disvalue so that some (large) number of people having their arm temporarily broken is worse than a single individual suffering from one hour of agonizing pain (a very severe harm in *the transmitter case*). However, it is inappropriate to allocate infinite moral disvalue to the harm of someone breaking their arm as it seems that there is some finite number of minor harms (e.g. the harm caused by being prevented from watching a World Cup game) in such a way that it is worse to cause the large number of minor harms instead of causing the single moderate harm. As such, ascribing infinite moral disvalue to some limited class of severe harms would result in the implausible result of many moderate harms never being able to axiologically outweigh these severe harms. Furthermore, for the reasons expressed in the discussion of the other horn of the dilemma, it is difficult to motivate and defend a specific qualitative level of severity at which harms change from infinite badness to finite badness. Whatever level was selected, there would be some number of harms slightly less severe than the least severe harm of infinite disvalue that would be intuitively worse than one instance of the least severe harm of infinite disvalue. For these reasons, we ought to reject the view that claims some harms are of infinite disvalue.

There are other ways in which someone can deny (2). Mayerfeld suggests denying that minor harms such as hangnails are of finite moral disvalue: one could claim they are not bad at all because they seem to be of negligible moral badness[110]. However, that the badness is negligibly small does not mean it is of no moral disvalue. If we had the choice of preventing the pain caused by a hangnail at no cost to us or to do nothing, the intuitively morally right thing to do is to prevent the pain because of that pain's disvalue. Similarly, it seems that five hundred people suffering from hangnails is worse than a single person suffering from a hangnail, and given the option, we would prefer to prevent the larger group rather than the single individual from that harm. A hangnail's pain does seem to be of some small moral badness in such cases. This denial of (2) becomes even less plausible when the qualitative severity of the minor harm being considered increases. In Scanlon's original *transmitter case*, some individuals (such as soccer fanatics) watching the World Cup game really enjoy the match and therefore the harm caused to them by preventing them from watching the game live is far greater than a hangnail. It would be implausible to claim that the harms to them are not of any disvalue at all. The value to the audience is clear. In many cases such fans expend significant effort and money to watch the televised match at special locations and both governments and businesses spend money to put on special televised events.

The above criticisms motivate a more subtle version of this denial of (2). On such a view, the minor harms that the Second Standard View applies to are only

[110] Mayerfeld, *Suffering and Moral Responsibility*, 187.

of no value in choice contexts when disparately larger harms are at stake. Analogously, minor benefits are not of moral value when much larger benefits are at stake. The harms *are* of minor disvalue when larger harms are not options in one's choice context. It is thus a context sensitive view. Call this view the *no value contextualism*. This view is *similar* to both Kamm's claim that such minor harms are irrelevant utilities (roughly, utilities that do not provide reasons) when compared with more severe harms and Scanlon's claim that such minor harms are not morally relevant when the more severe harms are at stake. The denial of (2) under present consideration is distinct from both Kamm's and Scanlon's views. This is because while both Scanlon and Kamm vindicate the deontic version of our intuitions in disparate conflict cases, they do not defend the axiological Second Standard View.

No value contextualism orders the 'worse than' relation in the following way: a situation with m' *major harms* and n' *minor harms* (and nothing else relevant to the situation's value) is worse than a situation with m *major harms* and n *minor harms* (and nothing else relevant to the situation's value) if and only if either m' is greater than m or m' and m are both zero and n' is greater than n. The view orders the 'equally valuable as' relation in the following way: a situation with m' high-value harms and n' low-value harms (and nothing else relevant to the situation's value) is equally good as a situation with m high-value harms and n low-value harms (and nothing else relevant to the situation's value) if and only if m'=m and m is greater than zero, or m'=m=0 and n'=n.

I have two objections to *no value contextualism*. The first objection is that the view makes the value of a benefit or harm implausibly depend on other options in the choice context. So whether a particular harm is of disvalue depends not only on intrinsic features of the harm itself but the other options available. This may initially appear surprising, but value often exhibits interaction effects like this where the value of one option depends crucially on the choice context. To use another restaurant example, wasabi may be of limited value when consumed on its own: however, used appropriately it can complement sushi and sashimi so that sushi consumed with wasabi is better than the value of consuming wasabi and sushi separately. Furthermore, because of the way drinks can complement food, a good wine list in a restaurant with a good menu is more valuable than the same wine list at a restaurant with a bad menu. The same wine options are more valuable to the diner given the context of her choices. In these cases, the value of an option depends on the other choices available. It is therefore not a good objection to *no value contextualism* that the value of options does not depend on other options in an agent's choice context.

The second objection to *no value contextualism* is more significant. According to *no value contextualism*, some minor harms are of no disvalue when in the same choice context as sufficiently more severe harms. This has implausible implications. Such a view would claim that given a choice context where we can either save one innocent person from agonizing torture or another innocent person from equally agonizing torture and cure ten people from mild short-lived headaches, both options are equally good. Kamm would argue that in such a

situation, the minor headaches are irrelevant utilities and so would not give us additional reason to select the option where in addition to preventing the torture we cure ten headaches. Unsurprisingly, consequentialists have objected to this deontic claim as implausible. The axiological version of the claim is more implausible than Kamm's deontic claim. It claims that the second option is not worse than the first. To bring out the implausibility further, change the case so you can save the same person from agonizing torture in both cases, but the only difference is that you can also cure ten other individuals' headaches at the same time. *No value contextualism* claims that both options are equally good - it is not better to cure the ten headaches at the same time! While no value contextualism is more plausible than claiming minor harms are of no disvalue in all contexts, it still has implausible implications. Minor harms must axiologically count for something, even in choice contexts where much more severe harms are at stake.

In order to respond to this second compelling objection against *no value contextualism*, one ought to instead claim that in choice contexts where sufficiently severe major harms are at stake, all sufficiently minor harms retain some, albeit limited, axiological significance. In such contexts, when comparing the value of two outcomes, it is worse if more minor harms are suffered rather than fewer and each minor harm is of infinitesimal disvalue. Call this view *tie breaking value contextualism.* While both views are contextualist, the views are significantly different. According to *no value contextualism*, the minor harms in such contexts are not of any disvalue and so minor harms never affect the

axiological ordering in such cases. However, according to *tie breaking value contextualism*, while minor harms never axiologically outweigh major harms, they do affect the axiological ordering. According to *no value contextualism*, the 'worse than' relation is ordered as follows: a situation with m' *major harms* and n' *minor harms* (and nothing else relevant to the situation's value) is worse than a situation with m *major harms* and n *minor harms* (and nothing else relevant to the situation's value) if and only if either m' is greater than m, or m' is equal to m and n' is greater than n.

Returning to the case from two paragraphs ago, unlike *no value contextualism, tie breaking contextualism* is able to vindicate and explain our moral judgment that preventing the torture of one person as well as the headaches of ten people is better than just preventing the torture of one person. It also explains why causing two instances of torture is worse than any number of minor headaches and one torture. This is because no number of infinitesimally disvaluable minor harms ever axiologically outweigh a severe harm. As a result, *tie breaking value contextualism* is not subject to the second objection raised against *no value contextualism* as minor harms are always of some disvalue, no matter what the context. Furthermore, this view also enables us to vindicate the Second Standard View by rejecting (2): it explains why no number minor harms to *the many* ever axiologically outweighs sufficiently severe harms to *the few* by claiming that each of the minor harms are of only infinitesimal disvalue.

In addition, as *tie breaking contextualism* does not deny (1), it respects a particularly plausible form of axiological impartiality by ensuring that within a given choice context that each and every harm of equal severity, no matter who suffers them, makes an equal contribution to the axiological evaluation of an outcome. Unlike Mayerfeld's view invoking *non-contribution point* or Carlson's diminishing disvalue of harms, tie breaking contextualism ensures that all qualitatively equal harms count for the same within a context. This is another substantial virtue of the view.

Conclusion

There are a number of widely held, defensible moral judgments that support Temkin's Second Standard View. However, as I have argued, the Second Standard View is incompatible with two particularly plausible and widely held axiological views[111]. This paper evaluated a number of attempts to resolve this incompatibility from the philosophical literature and found them all to be implausible. Instead, I proposed and defended a new view about how to resolve the incompatability, *tie breaking value contextualism*, that avoids important objections while capturing two sets of important intuitive judgments - that minor benefits to *the many* ought never to axiologically outweigh major benefits to *the few* and that each and every benefit and harm should contribute

[111] The Second Standard View, when combined with some other plausible axiological views, is also inconsistent with the transitivity of the better than relation. This inconsistency is argued for in Temkin, *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning,* 26-66.. In my paper 'Making Betterness Behave', I argue that if 'better than' is non-transitive, the 'more reason to bring about' relation should not be co-extensive with the 'better than' relation.

equally to the axiological evaluation of a state of affairs. While that view involves the sacrifice of an intuitively plausible axiological principle that claims that all harms are of finite disvalue, because of the incompatibility between the sets of axiological views considered, no view on this topic is without intuitive costs. Furthermore, *tie breaking value contextualism* offers the least intuitive sacrifices of all views considered and thus it offers a reflective equilibrium endorsed justification of the value of benefiting *the few* over *the many*.

Bibliography

Summary Measures of Population Health: Concepts, Ethics, Measurement and Applications. Edited by Christopher Murray, Joshua Salomon, Colin Mathers et al., . Geneva, Switzerland: World Health Organization, 2002.

"FIFA Media Release: Almost Half the World Tuned in at Home to Watch 2010 FIFA World Cup South Africa." [cited 2012]. Available from http://www.fifa.com/worldcup/archive/southafrica2010/organisation/media/newsid=1473143/index.html.

Bader, Ralf. "Neutrality and Conditional Goodness." (Unpublished Manuscript).

Bentham, Jeremy. *An Introduction to the Principles of Morals and Legislation (Collected Works of Jeremy Bentham).* Clarendon Press, 1996.

Broome, John. *Weighing Lives.* Oxford: Oxford University Press, 2004.

Carlson, Erik. "Aggregating Harms - should we Kill to Avoid Headaches?" *Theoria* 66, no. 3 (2000): 246-255.

Kamm, Frances Myra. *Intricate Ethics : Rights, Responsibilities and Permissible Harm.* New York, N.Y. ; Oxford: Oxford University Press, 2007.

Mayerfeld, Jamie. *Suffering and Moral Responsibility.* New York ; Oxford: Oxford University Press, 1999.

Mill, John Stuart. *On Liberty.* Boston: Ticknor and Fields, 1863.

Norcross, Alastair. "Great Harms from Small Benefits Grow: How Death can be Outweighed by Headaches." *Analysis* 58, no. 2 (1998): 152-158.

Otsuka, Michael. "Saving Lives, Moral Theory, and the Claims of Individuals." *Philosophy & Public Affairs* 34, no. 2 (2006): 109-135.

Parfit, Derek. *On what Matters.* Oxford University Press, 2011.

Scanlon, Thomas. *What we Owe to each Other.* Cambridge, Mass. ; London: Belknap Press of Harvard University Press, 1998.

Taurek, John. "Should the Numbers Count?" *Philosophy and Public Affairs* 6, no. 4 (1977): 293-316.

Temkin, Larry. *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning.* Oxford University Press, 2012.

Making Betterness Behave - Constructing a Transitive Deontic Ordering from a Non-Transitive Axiological Ordering [112]

A number of philosophers have argued that the 'better than' relation (hereafter >) is not transitive[113]. That is, it is not true that for all A, B and C, that if A>B and B>C, then A> C. Such a view is claimed to have many theoretical costs including threatening our conception of practical and moral reasoning, undermining versions of expected utility theory, making practical dilemmas impossible to avoid, and causing agents to be susceptible to *money pumping*. This paper will argue that many of these supposed costs can be avoided or, at least, largely mitigated. These costs mainly arise because most philosophers posit a particular connection between the axiological and the deontic. The dominant view is that both the axiological and deontic relations, > and 'more reason to bring about rather than'[114] are both transitive. Others, such as Larry Temkin, have suggested that both of these relations are not transitive. However, I will

[112] I am grateful to Ralf Bader, Nick Beckstead, John Broome, Tim Campbell, Will Crouch, Johann Frick, Felix Pinkert, Julian Savulescu, Guy Sela & Oxford University's Applied Ethics seminar & New York University's Axiology Working Group for comments on an earlier version of this paper.

[113] See, for example Larry S. Temkin, "A Continuum Argument for Intransitivity," *Philosophy and Public Affairs* 25, no. 3 (1996): 175-210. Another such argument is available in Stuart Rachels, "Counterexamples to the Transitivity of Better Than," *Australasian Journal of Philosophy* 76, no. 1 (1998): 71-83. Other arguments for non-transitivity can be constructed from cases not originally intended to serve that purpose, such as Parfit's Repugnant Conclusion (see Section 1 of this paper).

[114] In this paper, I am assuming that in deciding what we ought to do morally, the appropriate evaluand are the different states of affairs brought about by each possible action. This state of affairs includes the action itself.

argue that there is an important part of logical space that will capture many of the intuitions that motivate philosophers to affirm the non-transitivity of > while avoiding the aforementioned costs. In particular, I contend that if > is not transitive, one ought to take the position that 'more reason to bring about rather than' *is* transitive. Call this the *conditional non-coextensive thesis*. The focus of this paper is not to defend the conditional's antecedent, but to show that it is possible to generate a complete transitive deontic ordering out of a non-transitive axiological ordering on the basis of intuitively supported principles. To this end, the paper is structured as follows. Section I presents some of the arguments for the claim that > is not transitive and also illustrates some of the ways in which non-transitivity of > can occur. Section II outlines the practical and theoretical costs associated with accepting that better than is a non-transitive relation. I will then argue that these costs can be largely avoided. Section III outlines two independently motivated deontic principles that generate transitive 'more reason to do rather than' orderings on the basis of non-transitive 'better than' orderings. There are two kinds of deontic orderings I will propose and defend - one that applies when cardinal axiological comparisons are possible and one that applies when ordinal, not cardinal, axiological comparisons are possible. Section IV will illustrate the theoretical and practical advantages of my proposed deontic orderings. Section V will defend my two deontic principles from three important objections.

Section I - Arguments for the Non-transitivity of Better Than

There are many arguments in the philosophical literature for non-transitivity of the better than relation. I will not present or evaluate all of them here, but merely present two such arguments. The presentation will serve two purposes – it will motivate the problem as well as illustrate two different ways in which the better than relation may be non-transitive. This is important because as I will argue that the notion of non-transitivity is insufficiently fine grained. Some objections that have been presented against non-transitivity of better than only apply to certain types of non-transitivity.

I will present one argument based on intrapersonal benefits and burdens from Larry Temkin, which I have selected on the basis of its simplicity[115]. His argument is based on three claims:

> Claim 1: For any unpleasant or 'negative' experience, no matter what the intensity and duration of that experience, it would be better to have that experience than one that was only a *little* less noticeably qualitatively intense but twice as long.
>
> Claim 2: There is (or could be) a range of experiences, consisting of a *finite* number of unpleasant or 'negative' experiences ranging in qualitative

[115] This example was originally inspired by a 1993 unpublished undergraduate thesis by Stuart Rachels, but this formulation is based on the one used in Larry Temkin, *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning,* Oxford University Press, 2012), 135. There is also an analogous interpersonal version of this argument in ibid., 26-66.

intensity from extreme forms of torture to the mild discomfort of a hangnail, such that one could move from the harsh end of the range to the mild end in a finite series of steps, where each step would involve the transformation from one negative experience to another that was only a *little* less intense than the previous one (transformations of the kind in Claim 1).

Claim 3: A mild discomfort for the duration of one's life would be better than two years of excruciating torture, no matter the length of one's life.

Now consider two very lengthy lives that you can choose between, L_1 and L_2. L_1 and L_2 are exactly analogous in all relevant respects except that L_1 contains two years of excruciating torture and L_2 contains four years of torture whose intensity is almost as bad as L_1's, but is clearly and noticeably less bad[116]. In accordance with Claim 1, $L_1 > L_2$. Next compare, L_3 with L_2, where L_3 stands to L_2 as L_2 stands to L_1. Again, in accordance with Claim 1, L_2's four years of intense pain seems better than L_3's somewhat less painful eight years of pain. One can continue these pairwise comparisons between L_3&L_4, L_4&L_5, L_5&L_6 and so on and by analogous reasoning, the first of the two lives being compared is better than the second. Eventually, in accordance with Claim 2, one would be comparing two alternatives with extremely minor harms, say, L_{20} and L_{21} – L_{20} suffering from an annoying hangnail for a very long time (many thousands of years) and L_{21} suffering from a noticeably milder hangnail that was twice as long as L_{20}'s. One thus has a sequence where $L_1 > L_2, L_2 > L_3, L_3 > L_4, L_4 > L_5$ ..., $L_{19} > L_{20}, L_{20} > L_{21}$.

[116] If such pains can be effectively quantified, you could imagine that L_2's pain is 90% of L_1's.

If > is transitive, then $L_1 > L_{21}$. However, Claim 3 says that L_1 is not > L_{21} – no matter how long one lived, the real but mild discomfort of a hangnail throughout one's life is preferable to two years of excruciating torture. This claim implies $L_{21} > L_1$ and therefore that > is not transitive[117].

Arguments of this kind imply a particular kind of non-transitivity, namely that the better than relation is *cyclic*[118]. Formally, the better than relation is cyclic if and only if there exists a finite sequence of distinct elements $x_1, x_2, \ldots, x_k$ (such as the alternatives in the above L series) in such a way that k is greater than 1 and x_j is better than x_{j+1} for all $j \in \{1,2,\ldots,k\text{-}1\}$ and x_k is better than x_1. Cyclicity of better than is a sufficient but not necessary condition for non-transitivity of better than.

---

[117] There is an analogous interpersonal version of this argument that depends on two claims: (C1) The benefit of preventing one quadriplegia is better than the benefit of preventing any number of one minute, mild headaches. (C2) There is a finite sequence of benefits, ordered in decreasing quality, $B_1, B_2, \ldots, B_j$, where B_1 is preventing one quadriplegia and B_j is preventing a one minute mild headache, such that for all f $\in \{1,2,\ldots,j\text{-}1\}$, there is some finite number of B_{f+1} that is better than any finite number of B_f. (C1) is an instance of what Temkin calls the *Second Standard View.* I defend both C1 and a variant of Temkin's Second Standard View in my paper 'Valuing the Few over the Many'. If (C2) and transitivity of the better than relation is true, that entails the denial of (C1) as preventing some finite number of one minute mild headaches better than preventing one quadriplegia. If (C1) and (C2) are true, then the better than relation will be non-transitive and cyclic.

[118] Another argument for cyclicity and non-transitivity of the better than relation can be constructed from Derek Parfit's Repugnant Conclusion. To see this, consider:

> *Parfit's Repugnant Conclusion*: Consider a series of population alternatives, where the first in the series is 100 people all enjoying an extremely high quality of life, and each subsequent alternative in the series has twice as many different people all enjoying slightly but noticeably less good lives that are still worth living. Within each option, all individuals experience the same quality of life. The last alternative in the series is a large group of billions of people living a life barely worth living, in Parfit's description they are living listening to just muzak and eating just potatoes.

Each subsequent alternative in the series is intuitively better than the last because there is a large net increase in total wellbeing, however the last member of the series is intuitively worse than the first option in the series. This case lends further intuitive support to the non-transitivity of >. Derek Parfit, *Reasons and Persons,* (Oxford: Clarendon Press, 1984), 384-90.

There are arguments that, if sound, establish the non-transitivity of better than but not the cyclicity of better than[119].

It is also possible to claim that better than is non-transitive but also non-cyclic. For the purposes of this paper, it is sufficient to show how a non-cyclic non-transitive view about 'better than' could arise. For example, consider the view that variable size population axiological comparisons are only possible if the size of the populations are different by one person. So if population K1 consists of 100 people, K2 consists of 101 people and K3 had 102 people, then assuming K1> K2, K2> K3, then there would be a case of non-transitivity due to non-comparability between K1 and K3. The distinction between cyclic and non-cyclic types of non-transitivity is important, as I will argue that many of the objections to non-transitivity of better than only apply to cases involving cyclicity.

I am not concerned to defend the non-transitivity of the better than relation in this paper. There may be other ways to account for the cases that suggest better than is non-transitive and I will not evaluate them here. Instead, the primary purpose of this paper is to argue that the denial of the transitivity of > does not need to have many of the philosophical costs often associated with embracing the view. It is important to know the merits and disadvantages not only of one's own philosophical views, but also opposing views.

[119] Ralf Bader presents and defends a view where a number of different 'better than' relations are non-transitive and non-cyclic in Ralf Bader, "Neutrality and Conditional Goodness," (Unpublished Manuscript). For the current purpose of distinguishing cyclic and non-cyclic versions of non-transitivity, it is not necessary to give the details of his axiological views.

Section II - Supposed Costs of Non-transitivity of Better Than

There are five theoretical problems that the non-transitivity of the better than relation is claimed to have. It (i) leaves agents open to being *money pumped*, (ii) leads to *practical dilemmas*, (iii) removes the justification for the use of pairwise comparison to eliminate options in practical reasoning, (iv) denies an axiom of some variants of expected utility theory and (v) that the meaning of word better than necessarily implies that it is transitive. I will elaborate the first three problems in this section. I will not discuss (iv) and (v) as there is already an excellent discussion in Larry Temkin's *Rethinking the Good*. I will then turn to responses to these problems in the next section.

II.A - Money Pumps

The first two problems do not apply to all cases where the better than relation is non-transitive, but the *cyclic* subset of such cases. One example would be A>B, B>C & C>A. Or to take Temkin's example from Section I,

$L_1>L_2$,$L_2>L_3$,$L_3>L_4$,$L_4>L_5$ ..., $L_{19}>L_{20}$,$L_{20}>L_{21}$ but also $L_{21}>L_1$.

An agent acting according to such a set of cyclic better than comparisons could be *money pumped* in the following way. For simplicity, consider the aforementioned cycle A>B, B>C & C>A. Firstly, imagine an agent, call her Kelly, is in a situation in which B will be brought about if she does nothing or she can

bring about A for the cost of ten pence. When given this option to bring about A instead of B (in the absence of any other options, such as C), all other things being equal Kelly should do so as the axiological improvement is more valuable than ten pence[120]. Imagine then Kelly is given the option to bring about C (instead of A) for ten pence (option B is no longer available). Acting again to bring about the better option, all other things being equal, Kelly should pay the small amount. However, following analogous reasoning, when subsequently given the option to pay ten pence to bring about B instead of C (now that option A is currently unavailable), Kelly would do so, resulting in the situation that Kelly started with while leaving Kelly thirty pence poorer. More precisely, the cyclic axiological comparisons that guide Kelly are A-10p>B, C-20p>A-10p & B-30p>C-20p. Furthermore, if she stubbornly keeps her cyclical preferences, then Kelly can continue making these small payments to change outcomes until she has spent all of her money. It is absurd that a theory of rational action would require us to act according to such cyclic preferences.

II.B - The Practical Dilemmas of Never Acting Rightly

According to Parfit[121] and Norcross[122], it is impossible to not act wrongly when better than *cycles* between a number of options. This conclusion can be reached by assuming that an action is wrong if there is another alternative that

[120] If you disagree that the moral value is equal to this small amount of money, then reduce the amount appropriately – perhaps one Japanese yen?

[121] Derek Parfit, "Justifiability to Each Person," *Ratio: An International Journal of Analytic Philosophy* 16, no. 4 (2003): 368-390.

[122] Alastair Norcross, "Contractualism and Aggregation," *Social Theory and Practice* 28, no. 2 (2002): 303-314.

you can bring about that is better[123]. After all, in the aforementioned simple cyclic case if you chose A, C>A; if you chose B, A>B; and if you chose C, B>C. That is, for each option available (A,B or C), there is always another option that is both morally preferable (i.e. one ought to do instead) and better than that option when making pairwise comparisons between two options at a time. This would lead to moral dilemmas where we could never act rightly.

Furthermore, these cyclic moral dilemmas are unlike more standard moral dilemmas where all of the agent's actual options are very bad outcomes. One such case is from *Sophie's Choice*, where Sophie is told by a guard in a Nazi concentration camp that she must select one of her children to be killed by the guard otherwise the guard will kill both[124]. All of her options are bad and some philosophers think that all actions are wrong, making this a genuine moral dilemma. What is striking about the purported moral dilemmas caused by non-transitivity of better than is that they can occur even when all of the options available are *good* alternatives! If this really is an implication of affirming non-transitivity of better than, then this gives us another reason to reject the claim that better than is non-transitive. This is distinct from the previous money pump problem raised in the previous subsection. The money pump problem arises when there is a sequence of different choice contexts, whereas this problem arises within a single choice context.

[123] That is, any action Φ is wrong if one can bring about a better outcome by doing not- Φ. This is a standard claim about wrongness often made by consequentialists.
[124] William Styron, *Sophie's choice,* Random House, (1998)

II.C - Removing the Justification of Pairwise Comparison

Transitivity of the better than relation justifies one of the ways we choose between a large number of possible choices in everyday practical decision making. Temkin argues that if you affirm non-transitivity of better than, this important kind of reasoning loses this justification[125]. When faced with any complex decision with multiple alternatives, assuming transitivity of better than, we can use a number of pairwise comparisons that will lead to the best of all options. We can focus on just two options at a time using pairwise comparison and eliminate the less good option from further consideration. The *winner* of the previous comparison (i.e. the better of the two) is then compared with another previously unconsidered option, repeating this procedure until the best option is left. So from any set of *n* options, whether they are products, graduate schools or jobs, when each pairwise betterness comparison is clear, one can choose the best from *n-1* comparisons, each one having one *winner* and *loser* (the worse of the two options). The final winner is the best option, which assuming the transitivity of better than, is better than all other options. However this clear, systematic approach that can be relied upon in everyday practical decision making presupposes transitivity of the better than relation. There is no reason to justifiably remove a loser from further comparison if that option may turn out to be better than one of the later winners. What justifies

[125] See Chapter 13 of Temkin, *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning,*. It should be noted that this argument is about the *justification* of one particular kind of practical reasoning and it is irrelevant whether or not actual agents often use this kind of reasoning. Not everyone uses this kind of practical reasoning. Actual agents use various heuristics and types of reasoning.

removing a loser from further comparison is that the final winner will be better than each and every loser, either directly via a pairwise comparison, or better than an option that is in itself better than that loser. However, if transitivity of the better than relation is false, then it could be the case that the final winner of such a procedure is in fact worse than one of the previous losers. Call this the *knocking out* method.

If non-transitivity of better than has all of these costs and it is evident in both interpersonal and intrapersonal comparisons, then affirming non-transitivity of the better than relation is grave indeed. In many public health decisions, such as deciding which treatments to fund, there would be no right option. Acts of everyday practical reasoning, from picking which pieces of clothing to wear in the morning to choosing food from a menu to deciding whom to give a scholarship all would all rest on a faulty justification. These costs would wreak havoc throughout our moral lives and be devastating to much of applied ethics. Therefore we have good reason to see if these costs are as bad as they are presented and if so, to reconsider the non-transitivity of better than.

Section III - Mitigating the Costs

III.A - Motivating The Difference between the Axiological and the Deontic

It should be noted that all three problems stated in Section II are problems of *how to act* supposedly stemming from a surprising feature of *value* (namely the non-transitivity of the better than relation). It is not surprising, then, that those

that are most concerned with the problems raised by non-transitivity of the better than relation are philosophers that posit a tight connection between the axiological and the deontic such as certain consequentialists. Maximizing consequentialists claim that the only right acts are those that best promote value, so that any act Φ is wrong if one can bring about a better outcome by doing not Φ. However it is not only consequentialists that have views that entail the practical problems described in Section II. Certain buck-passers about value hold the view that claims about value just reduce to claims about reasons, so when an option A is better than another option B, one has more reason to bring A about than B. If such a buck-passer claimed that any action Φ is wrong if one has more reason to do not Φ, then given non-transitivity of better than, they would be subject to the practical dilemmas described in Section II.B.

This problem for certain consequentialists and buck-passers gives us an idea of how to circumvent these practical problems, namely by positing a different connection between the axiological and the deontic. Non-transitivity of > seems to threaten our practical reasoning because the pairwise better than relation is often assumed to be co-extensive with the 'more reason to bring about' relation[126]. If > and the more reason to bring about rather than relation are co-extensive, then if > is cyclic and non-transitive, the more reason to bring about relation is also cyclic and non-transitive, and this would threaten practical rationality by leading to the irrational cyclic preferences and inability to select one option as most choice worthy in the cyclic cases mentioned previously.

[126] That is, it is often assumed that if state of affairs SOA_1 is better than another state of affairs SOA_2, if given the option we morally ought to bring about SOA_2 instead of SOA_1.

However, I will argue that if > is non-transitive, > need *not* be co-extensive with the more reason to bring about relation. As I will argue, this *conditional non-coextensiveness thesis* allows us to avoid many of the aforementioned problems.

There are many intuitive differences between the axiological and the deontic, which one need only look at the rich debate between consequentialists and nonconsequentialists to find[127]. There are also intuitions that are specific to the *non-coextensiveness thesis*. We have no intuitive reasons to believe that a formal property such as transitivity is held by the complex, potentially queer, world of value in all contexts. It would be theoretically neat and useful if it was so, but I have no intuitive presumption that better than is transitive. Furthermore, as Larry Temkin has shown, there are numerous widely held intuitions that taken together are inconsistent with transitivity of the better than relation. So we have some intuitive reason to believe that better than is non-transitive. The deontic is not analogous to the evaluative in this way. Unlike in the axiological case, we do have reasons, intuitive and otherwise, to believe that the formal property of transitivity of the 'more reason to bring about rather than' relation ought to be held by a theory of rational or moral action. This is precisely for the objections raised in Section II. It is intuitively absurd that a normative theory would claim we rationally or morally ought to be voluntarily money pumped or that there is no right action in the cases Parfit and Norcross discuss.

[127] For example, see David Ross, *The right and the good,* Oxford University Press, USA, 2003); Frances Myra Kamm, *Intricate ethics : rights, responsibilities and permissible harm,* (New York, N.Y. ; Oxford: Oxford University Press, 2007), 509.; Robert Nozick, *Anarchy, state, and utopia,* Basic Books, 1977).

I will first consider but ultimately reject a deontic principle that makes prescriptions about how agents should act when presented with non-transitivity of better than over a sequence of choice contexts. This principle both prevents one from being money pumped and provides guidance about how to act in sequences like Temkin's L series. Even though I will ultimately reject this principle, it is instructive to consider it.

III.B - The Commitment Deontic Principle & Kamm's Axiological Prioritization

Consider the following deontic principle that will avoid many of the purported problems of non-transitivity of better than.

> *The Commitment Principle* states that all other things being equal (that is, when the only differences between options is axiological), one has more reason to bring about a new state of affairs over a current state of affairs only when the new state of affairs is *both* better than the current state of affairs and other *relevant* previously available states of affairs.

This principle requires further explication. It is best explained by example. In the following cases, unless I mention otherwise, I will assume that the only morally relevant differences between the options are axiological. Take the aforementioned simple cycle, A>B, B>C, C>A. When presented with the first

possible trade between A and B, according to the *Commitment Principle* she morally ought to pay ten pence to bring about A instead of B when there are the only two options available (thereby justifiably foregoing option B). However, when option C becomes available, she should *not* pay a further ten pence to bring about C instead of A-10p (even though C-20p>A-10p) because C-20p is not morally superior to B, a *relevant* previously available option. A previously available option is *relevant* for the purposes of the principle when it was justifiably foregone by an agent acting in accordance with the *Commitment Principle*.

This does not mean, however, that no new options can change one's original decision – that would be highly implausible. Remember that all other things being equal, you have more reason to bring about a new option when it is better than all currently available and previously justifiably foregone options. For example, now imagine that after the agent has paid ten pence to bring about A, a new option, D becomes available which is better than each of A, B & C. In such a case, she would be justified in bringing D about even for some monetary cost.

Furthermore, note that the principle does not state what one ought to do when one has a number of options that are better than both the current state of affairs and relevant previously foregone states of affairs. That is, it is neutral on whether one should be *maximizing* (i.e. it is only permissible to bring about the

best possible state of affairs) or *satisficing* (i.e. it is permissible to bring about any state of affairs that is better than some level of value).

In addition to avoiding money pumps, this procedure also provides an account of right action in long sequences such as Temkin's L series. On such a view, we can continue improving the situation by accepting each better than comparison *until* the next option is not better than *all* previously foregone options. Applied to *Parfit's Repugnant Conclusion*, we should stop the increases in the slightly worse off population when doing so brings about a worse state of affairs than the first group of 100. Considering the choice context during moral deliberation allows us to vindicate the non-transitivity of > while avoiding the 'money pump' problem of cyclic preferences and providing a plausible guide to action in difficult sequential cases like Temkin's L series.

The *commitment principle* has some plausibility. Agents do take into account their former decisions and all other things being equal, maintaining the status quo is often viewed as justified. However it is susceptible to two important objections.

Firstly, one implication of the commitment principle is that the mere order in which the options are presented will affect the deontic 'more reason to bring about rather than' ordering of the outcomes. Given the axiological cyclicity of the above case (A>B, B>C, C>A, A-10p>B, C>A-10p), if B and C are presented to the agent as options first, then, an agent following the *commitment principle*

will supposedly have more reason to stick with B than any other option even after being subsequently presented with A for the cost of ten pence. This is because the new option is worse than a previously justifiably foregone option (C). However, as shown above, if A-10p and B are the agent's first options, and then C is offered as an option later, the agent has more reason to select A-10p than any other option (as A-10p>B and B>C). Given that the agent is presented with all of options A-10p, B & C in both sequences, it is counter-intuitive that an agent will have a different deontic ordering of all options. Similarly, in Temkin's L series, whether one starts at L_{21} or somewhere else, say, L_{19} in the sequence may affect the deontic ordering of all options. This is because the option that is not better than L_{21} (the last member in the sequence) can be different to the option that is not better than L_{19}. To see this occur, let us assume that L_9 is the first presented option that is not better than L_{21}, whereas L_4 is the first presented option that is not better than L_{19}. If an agent acting according to the *commitment principle* was presented with a sequence, presented two options at a time, starting from L_{21} will continue down the series until L_{10} and not accept any further trades, even if presented with all of the remaining options from L_1 to L_{21}. According to the commitment principle, L_{10} is the option that the agent has more reason to bring about than any of the subsequently presented options. However an agent acting in accordance to *the commitment principle* would select L_5 if the sequence is started at L_{19}, even if all of the other options (including L_{20} & L_{21}) are made available later. The commitment principle gives a different deontic ordering of these two sets of options even though the options

available are precisely the same, they are just presented in a different order. Call this the *ordering* problem.

Secondly, the *commitment principle* does not provide guidance about how to act when all options are presented at once. As shown in the last paragraph, the deontic prescription provided by the principle depends crucially on the order in which the options are presented. However, when all options are available at once, there is no order to rely upon and I can see no reasonable, plausible way that is not ad hoc to privilege a particular order on which to base the *commitment principle* on. As such, it is at best only a solution when non-transitivity or cyclicity is instantiated across different sequentially presented options, as in the money pump problem. It does not provide a solution to the problem that Parfit and Norcross raise. Many benefit and burden distribution cases, like public health distribution decisions, provide many different options at once, rather than in a sequence. It is therefore important to have a deontic principle that can guide us in such cases. We ought to reject the *commitment principle* for these reasons.

Frances Kamm offers a deontic principle that can be used to decide how one should act even when there is non-transitivity of the better than relation and all options are presented at once[128]. She believes that one of the features that is axiologically relevant can be given priority during deliberation over which act we ought to do. In particular, she argues that qualitative differences in wellbeing

[128] Kamm, *Intricate ethics : rights, responsibilities and permissible harm,* 298, 485.

(i.e. how valuable something is for each person) should have priority over the quantitative differences (i.e. how many people benefit) in cases where there is a sufficiently different qualitative difference between the options being considered. According to this approach, with any case in which we are required to determine who to benefit/harm, one should consider the qualitative aspect not the quantitative aspect of the options first. That is, first treat the options as if there was only one individual being harmed (or benefited) in each possible option. If there are individuals suffering (or benefiting) from different qualitative levels of harms or benefits in the same option, then consider the strongest individual wellbeing claim from that group. All other things being equal, if the strongest opposing options are not sufficiently qualitatively similar, help the group[129] that has the qualitatively strongest moral claim (e.g. greatest individual harm prevented). If the strongest opposing options are sufficiently qualitatively similar, then help the group that contains the largest number of qualitatively relevant individuals.

This response gives differing action prescriptions depending on the different available options. This method is best illustrated via examples: imagine if we could prevent (I) the premature death of one innocent young person or (II) many permanent quadriplegias in innocent young people or (III) even more permanent paraplegias in innocent young people. According to Kamm, death is qualitatively similar but more serious than quadriplegia, paraplegia is qualitatively similar but less serious than quadriplegia, but death is not

[129] It could be a group of one.

qualitatively morally relevant to paraplegia. Call this case *disability/death prevention*. Given only options (I) and (II), Kamm would recommend option (II). Given all three options, Kamm would recommend option (II) because paraplegia is not qualitatively related to death but one can benefit more qualitatively relevant individuals by choosing option (II) and preventing many quadriplegias. Because paraplegia is not qualitatively morally related with death, Kamm would recommend (I) if (I) and (III) were available. On the other hand, if only (II) and (III) were possible, because someone's death was not at stake, Kamm's procedure would recommend (III). Her view is motivated by her belief that one ought not to prevent very minor harms to many people (such as curing the headaches of billions) instead of saving the life of one young innocent person.

A similar response giving priority of quality of benefits over duration of benefits could be applied to Temkin's L series. It would recommend choosing the longest member of the series where the burden that is sufficient qualitatively similar to the most severe burden, L_1. However Kamm's proposed solution relies on the prioritization of one or more axiological factors and I will show that some axiological orderings that instantiate non-transitivity of better than are not able to be plausibly prioritized. As a result, I will argue that we should reject Kamm's solution in favour of another view that does not require the prioritization of certain axiological factors.

Non-transitivity of the better than relation can be instantiated when different factors or the significance of those factors change when different comparisons are being made. Here is one difficult real life case where it is not plausible to believe that one such factor should take priority over the others at the deontic level. Some philosophy research job advertisements are listed as open rank. That is, they are open to applicants at all stages of their career, from graduate students finishing their PhDs to senior professors. In some such searches, assuming that they are genuinely open rank searches, sometimes each philosopher's merits for the job are evaluated in light of their stage of career, in such a way that different factors are relevant for different candidates. If they did not, then clearly established senior faculty applicants would often be awarded such jobs, as they have better research records, more teaching experience, and so on. For senior professors, both their record of research and reputation is important to evaluating their value as a candidate. They are expected to have a strong research record and international reputation. A PhD candidate, on the other hand, is not expected to have such a record or reputation, though they are expected to display research *potential*. For very senior candidates, research potential may not be important for their suitability as a candidate at all when compared to other professors. The department may accept that professor is not producing much research anymore, but can contribute to the department in other ways by attracting quality graduate students and increasing the reputation of the institution. The importance of

these factors on the evaluation of each candidate vary throughout the different stages of their career. I will explain how this may work by example.

For simplicity, let us stipulate that there are only four factors that can be relevant to the value of a job candidate - their research record, teaching experience, research potential and reputation. Assume each of these can be given a score from 0-10, where 0 is worst and 10 is best. Furthermore, assume that the units of different factors are comparable and can be traded off against one another in such a way that one point in one factor is as valuable as one in another, as long as that feature is relevant to the axiological comparison at hand. This is, of course, a gross simplification, but it is sufficient for illustrative purposes to show how such non-transitivity could be instantiated between options where multiple factors change in their importance depending on the different evaluations. Given the many *additional* considerations in actual decisions, it is even more likely that the non-transitivity of better than will be instantiated in a less simple real life case. I will consider a case where a cardinal ordering of the different features that contribute to each options' value is not available after this case. Assume 3 candidates, a senior professor (W), an associate professor (X), and a doctoral student who is about to graduate (Z). The respective scores on each dimension are displayed below and there are no other relevant differences between the candidates:

	Research Record	Teaching	Research Potential	Reputation
Senior Professor (W)	9	8	2	9
Associate Professor (X)	7	8	8	7
Graduate Student (Z)	1	6	9	1

Given the varying importance of the criteria, the search committee judge W as better than X. After all, W has a better research record (by two points), reputation (by two points) and equal teaching experience, and for professorial comparisons with W, research potential does not matter. It is good for the department to have someone who is suitably famous, can attract good graduate students, without very strong prospects of future research output. X is a better candidate than Z, as while research record & reputation does not matter for comparisons with graduate students, the associate professor has a better teaching record (by two points) and only one less point regarding research potential. However, when Z is compared with W, Z is a stronger candidate. In *that* comparison, research record and reputation are not relevant, and only teaching and research potential are relevant. Z has much better research potential (7 points) and only slightly worse teaching experience (2 points). Imagine, for example, that the student appears to her referees to be the next philosophical genius and has good teaching evaluations from all five years of their graduate program.

III.D - Constructing a Transitive Deontic Ordering

There are many possible ways of creating a transitive ordering out of a non-transitive ordering. The previously considered methods (the *commitment principle* and Kamm's axiological prioritization) achieved this by disregarding a subset of pairwise axiological comparisons. However, it is not clear in cases

such as the aforementioned job search that any single axiological factor (or subset of factors) could be given priority over all other factors and as such provide a principled basis on which to disregard certain pairwise axiological comparisons. How should we order the choiceworthiness of different options in such cases where all pairwise evaluative comparisons should be taken into account and the only differences between options are axiological? There is one straightforward way to take into account all of the axiological pairwise comparisons that is inspired by both voting theory and sporting competitions.

The type of cases is analogous to when we want to select the best participant or team in a competition where the competition involves many different skills. In actual sports competitions, all sorts of different methods are used for such purposes. However one method stands out because it can be used on the basis of pairwise comparisons only. In particular, I advocate using a method that takes into account the amount each candidate is better than each other candidate through a series of pairwise comparisons between all candidates[130].

[130] Larry Temkin considers a simpler axiological version of my method in Chapter 13 of Temkin, *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning,*. He calls it the *sports analogy model*. On such a model, how teams/outcomes axiologically compare all things considered, is a function of both how they compare to each other *and* other teams/outcomes like in a round-robin (also known as a all-play-all) tournament. On such a model, to work out which of two outcomes is better, one has to count the number of times it is better than (or beats, in the case of teams) each other team/outcome. The better team of the two compared is the one that wins more games and they are equally good if they win an equal number of games. My method is distinct for a number of reasons. Firstly, when it is possible to do so, like in the case currently under discussion, it takes into account not just whether or not one outcome is better than another, considered pairwise, but also how *much* one outcome is better than another. Secondly, my method is not intended to axiologically order outcomes, but to provide a transitive deontic ordering of the relation 'more reason to bring about than' from a non-transitive evaluative ordering. Because of these differences from the model Temkin discusses, it is not susceptible to many of the objections Temkin raises against it. I will discuss an objection that my method *is* subject to in the final section of this paper.

This is similar to a round robin tournament where each competitor competes with every other competitor in a series of pairwise competitions.

I will explain the method by example. Each candidate can sum the degree to which they are better ('wins') and worse ('losses') with each other option to get a choiceworthiness score[131]. W beats X by 4 points, but loses to Z by 5 points, for a total score of -1. X loses to W by 4 points, but wins against Z by one point, for a total score of -3. Z beats W by 5 points, and loses against X by one point, for a total score of 4. This score gives us a deontic ordering of 'more reason to bring about rather than' from highest to lowest score. As Z has the highest score, Z is selected according to this method that takes into account the amount that each option is better than another. This deontic ordering is transitive, which is evident as it can be represented on a real number line. When two options are axiologically non-comparable, pairwise comparisons do not contribute to either comparative's choiceworthiness score. Call this method the *cardinal choiceworthiness ordering*. The axiologically relevant comparisons are summarized in the following table:

Axiological Comparison	Axiologically Relevant Features	Contributions to Choiceworthiness Score
W & X	Research Record, Reputation, Teaching Record	+4 to W, -4 to X
X & Z	Teaching Record, Research Potential	+1 to W, -1 to Z
W & Z	Teaching Record, Research Potential	+5 to Z, -5 to W

[131] Options that are axiologically non-comparable with all other options will therefore have no wins or losses. This would result in a choiceworthiness score of zero, which would be higher than some options that lost more than they won.

It is also possible to use a similar method when we do not know (or there is not a fact of the matter) how much one option is better than another - that is, when we only have ordinal information about the betterness relations - or when the evaluative differences are cardinal but non-comparable between options[132]. Take the above case, but introduce a post-doc, Y and a second graduate student Z'. Assume that the betterness relations are as follows, without any information about how much one is better than another, so all we know is that W>X, W>Y, W<Z, W<Z' (W wins 2, loses 2), X>Y, X<Z, X ties with Z', (X wins 1, loses 2) Y>Z, Y>Z', (Y wins 2, loses 2), Z>Z' (Z wins 3, loses 1; Z' wins 2, loses 1). On the method I propose in cases where only ordinal axiological comparisons are possible, each candidate can be compared to each other candidate, and then ordered by the number of times they win such comparisons minus the number of times they lose. This number provides a deontic ordering of 'more reason to bring about rather than' from highest to lowest score. This is known as the *Copeland Method*[133]. The highest scoring candidate by this method is the most choice worthy option. Z gets the most wins over losses, and so given this method ought to be selected. Using this method, in some comparisons, there will be more than one most choice worthy option, as would be expected to be sometimes the case. This deontic ordering is also both complete and transitive. This is evident because every option has a sum of the number of wins over losses and this can be represented by a real number line. However, for clarity, I will focus on the cardinal version of the method unless otherwise specified.

[132] This method can also be used when the evaluative differences are cardinal but non-comparable between options.

[133] This method is ascribed to A.H. Copeland, who first presented this method at a graduate seminar at the University of Michigan in 1951.

Both of these methods generate an overall transitive deontic ordering of all the actual options out of pairwise axiological comparisons. The pairwise deontic relations of the actual options can then be derived from the overall transitive ordering. According to this way of generating a deontic ordering, there are no deontic pairwise comparisons that are independent of context because how choice worthy an option is depends on its axiological relations to other context determined options. As I will illustrate in more detail later, this means that the 'more reason to bring about rather than' relation between two options can change depend on the choice context and is in this way context sensitive.

Section IV - The Merits of My Transitive Deontic Orderings

IV.A - Theoretical Merits

I have presented a way of generating a transitive deontic ordering from a non-transitive axiological ordering. I will now consider the advantages and disadvantages of this ordering over others. This subsection of the paper will outline the theoretical advantages of this ordering. The next subsection will outline the practical advantages of the ordering, outlining how it avoids Parfit's and Norcross's moral dilemmas and avoids money pumps like those previous described in this paper.

Firstly, my two deontic ordering methods take into account all axiological comparisons rather than a subset of those comparisons. As a result the method does not require either a privileged order or prioritized axiological factor and is therefore not subject to the objections that the *commitment principle* and Kamm's prioritization succumb to. Secondly, where possible, it takes into account not only each axiological comparison but to what degree each option is better than each other option.

My proposed ordering is also importantly similar to the Borda Rule voting method[134]. That voting method ranks electoral candidates by assigning each candidate a certain number of points corresponding to the position in which that candidate is ranked by each voter. The higher a candidate is ranked, the more points they are assigned from that voter. The total number of points from all voters is then calculated for each candidate and the candidates are ranked according to those scores. On the most prominent variant of the Borda Rule, in an election with *n* candidates, a candidate will receive n points for a first preference vote, n-1 for a second preference vote, n-2 for a third preference vote, and so on until all preferences are counted.

My cardinal choiceworthiness method is similar in the following way. It treats each evaluative option as if they were a voter who ranks all of the options

[134] Thanks to Ralf Bader for highlighting this similarity. This similarity with the Borda Rule is important my method will share some of the same desirable properties as the Borda Rule over other possible ordering methods. For example, like the Borda Rule, my method will never rank a *Condorcet loser* as most choice worthy, where a *Condorcet loser* is an option that is worse in pairwise comparisons with all other options.

including 'themselves'. Call each such voter an *evaluative voter*. Each *evaluative voter* generates their ranking of all of the different evaluative options using their pairwise comparisons with each of the candidates. The number of points each evaluative candidate' gets from each evaluative 'voter' depends on how that candidate evaluatively compares to that voter (i.e. how much better or worse than the evaluative voter is when compared to each option) if such cardinal evaluative comparisons are possible. Each evaluative voter contributes nothing to his own score, because it is equally valuable to itself. This is the 'choiceworthiness' score that deontically ranks the options according to my method.

Using the original job market example, evaluative voter W contributes 5 points to Z and subtracts 4 points from X; evaluative voter X contributes 4 points to W and subtracts 1 point from Z; evaluative voter Z subtracts 5 points from W and contributes 1 point to X. This results in the same scores as my choiceworthiness method: Z has the highest score (4); followed by W (-1) and X (-3).

IV.B - Practical Merits of the Ordering

I have outlined some theoretical advantages of my proposed deontic orderings. I will now outline how they will solve the three practical problems outlined in Section II. Unlike maximizing consequentialist views, my deontic orderings can prescribe a right action when there are cycles in the better than relation. Of all

the options available, you have most reason to bring about the option that scores the highest choiceworthiness score by my methods. There are no options that an agent has more reason to bring about than those options, and so the moral dilemmas that Parfit and Norcross do not arise on my view.

Furthermore, in order to choose between options that instantiate non-transitivity of the better than relation, one can utilize my proposed deontic orderings to choose between options in a principled and intuitively supported way. My method succeeds where the *knockout method* fails. It succeeds even though both methods use the same pairwise axiological comparisons in order to deontically order the options. I will now consider how one can use my proposed deontic orderings to avoid being money pumped.

IV.C - Solving the Money Pump Problem

So far, my proposed deontic orderings only solve problems in which non-transitivity is instantiated within a single choice context and not across sequential choice contexts as in the money pump case. However my proposed solution can be extended to deal with sequential cases as well. In order to illustrate how my method can be used to solve a number of different variants of the money pump problem, I will set out different variants of the problem.

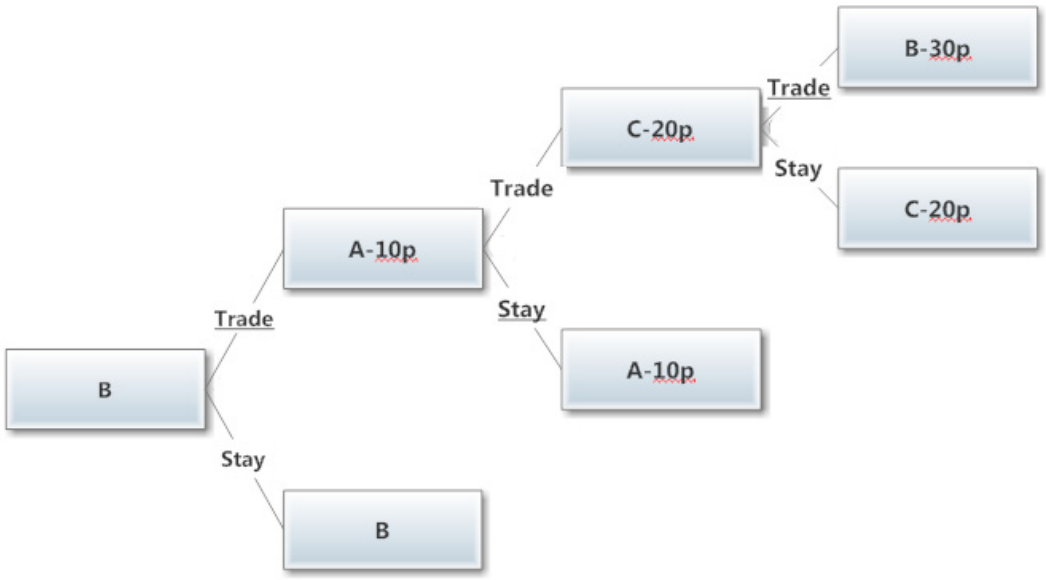
Money pump problems can be divided into cases where future options are foreseeable, and cases where future options are not foreseeable. I will outline an orthodox method of solving foreseeable money pumps and then argue that my proposed method is superior. The orthodox method of solving foreseeable money pumps involves an agent acting in accordance with *backwards induction*. While *backwards induction* solves a simple variant of foreseeable money pump problems, it is unable to avoid a more complex foreseeable variant of the case.

I will first outline the orthodox *backwards induction* method of solving the money problem and show that while that method solves a simple variant of the money pump problem, it is unable to avoid being money pumped in a more complex variant of the case. I will then show how my deontic ordering can be used to solve both the simpler and more complex variants of foreseeable money pumps. I will then proceed to show how my method can avoid money pumps that are not foreseeable.

The orthodox solution to money pumps when all future possible trades are foreseen is *backward induction*. According to this view, a chooser would not allow themselves to be money pumped because they would recognize that the first trade would lead to the second trade, the second trade would lead to the third trade, and so on, until they would be at their starting position minus the payments they had made. In particular, according to backwards induction the agent should set out a decision tree with all of the options available to them and first determine how they should rationally act at the last choice node of each

branch of the decision tree. On the basis of the rationality of each of those last node choices, the agent assumes that they would make that decision if they reached that node. The agent can then decide how they should rationally act at each of the immediately previous nodes on the basis of what outcome later in the decision tree that they would rationally pick. That is, they choose each option at the next to last node as if it would bring about their rational choices in the last node. They can then continue with this method from the last choice nodes to their first choice node to see what decisions are rational at each part of the decision tree. Using this method, the agent can thereby determine what trade would be rational at each node of the decision tree, relying on his future rationality, to predict what he ought to do at that node.

This is best explained by example. Consider again the aforementioned example of Kelly. She is faced with a case displaying cyclic better than relations, where A>B, B>C, C>A. Furthermore this cycle is still evident even when one makes ten pence payments as each step, so A-10p>B, C-20p>A-10p & B-30p>C-20p. For the purposes of the example, assume that B>B-30p, and B-30p is worse than any alternative better than B, so A-10p>B-30p. The agent is then offered three sequential trades as per the original example: to stick with B or trade for A-10p, then to stick with A-10p or trade for C-20p, then stick with C-20p or trade for B-30p. A decision tree setting out the agent's subsequent options is below:

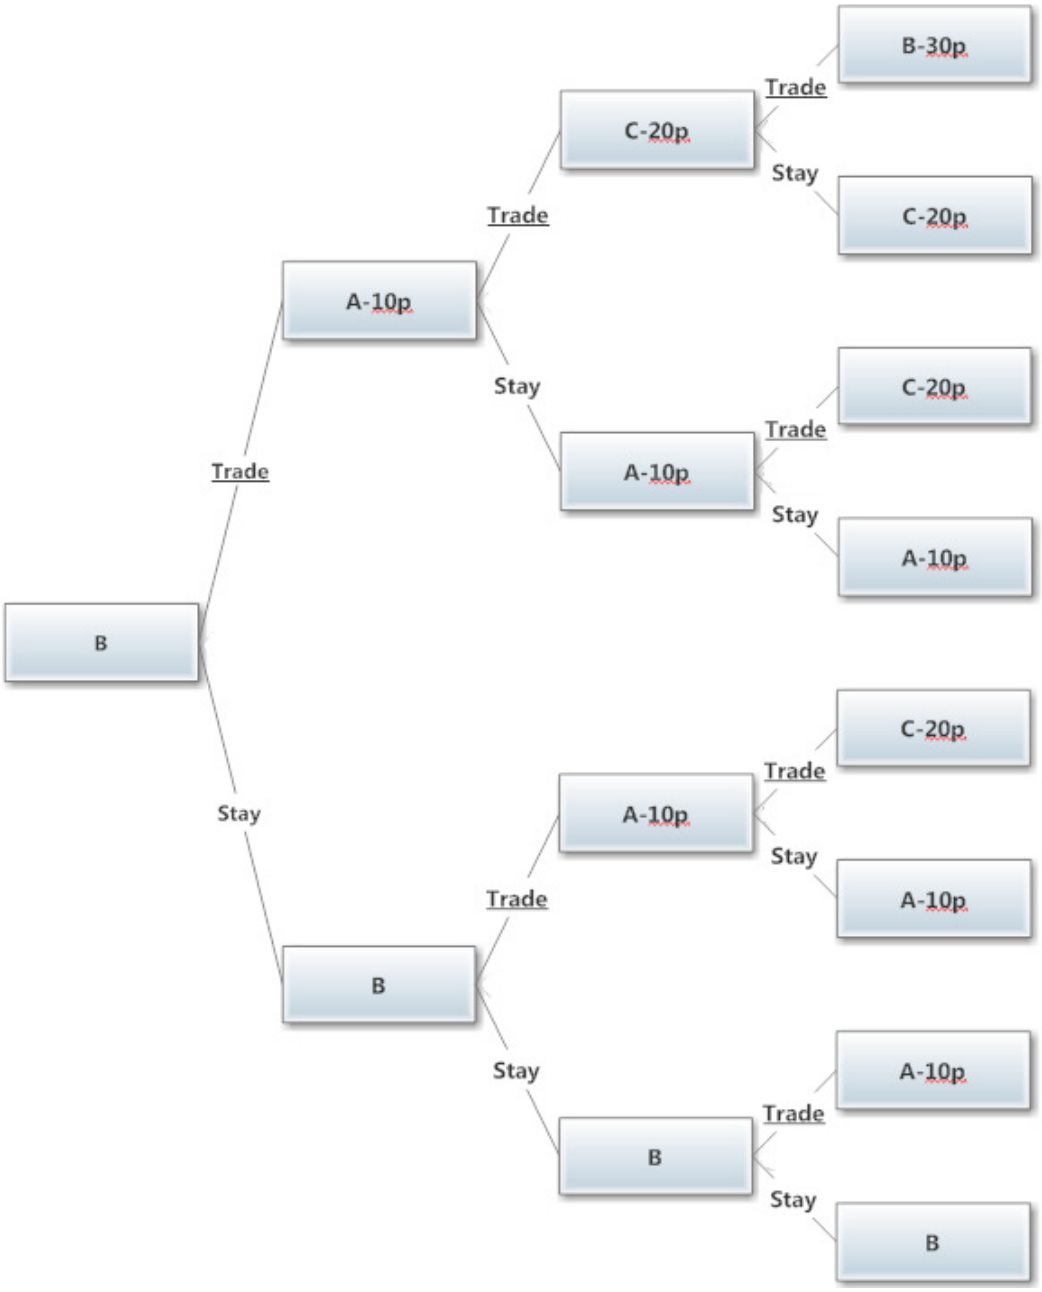


The underlined options indicate what an agent acting in accordance to backwards induction would rationally choose at each choice node. At the third and last choice node (pictured in the top right of the diagram), the agent would make the trade for B-30p, as B-30p>C-20p. At the second choice node, the agent following backwards induction can then decide how to act on the basis of how they would act on the last choice node, and so they act as if their choice is between B-30p (where they would ultimately end if they chose to trade at second to last node and acted rationally at the last node) and A-10p. However, as A-10p>B-30p, they would rationally choose to not make the trade and stick with A-10p. Then, knowing that at the second choice node they would select A-10p, at the first choice node, they would choose to make the trade to A-10p, as A-10p>B. The agent acting in accordance with backwards induction would choose to trade once, and avoid future trades, thereby avoiding being money pumped. In this three trade situation, they would never end up with the B-30p

outcome, which is the same situation they started in except with thirty less pence.

Backwards induction is an attractive solution to simple versions of the money pump problem. However, as Rabinowicz has shown, there is a more complex variant of the above case in which an agent acting in accordance with backwards induction will also be money pumped[135]. In particular, if the agent rejects a trade, the same offer will be presented again at the next stage and there are three stages at which offers are made. The axiological relations that guide the agent's action are the same as before - A-10p>B, C-20p>A-10p, B-30p>C-20p, B>B-30p and A-10p>B-30p. A decision tree setting out the agent's choices is set out below.

[135] Rabinowicz, W. 'Money Pump With Foresight' in Michael Almeida ed., *Imperceptible Harms and Benefits,* (Dordrecht, Holland: Kluwer Academic Publishers, 2000)



At every last choice node (i.e. the third trade opportunity, where the agent will have one of the following three trade options: trade for B-30p or stick with C-20p; trade for C-20p or stick with A-10p; trade for A-10p or stick with B), an agent will rationally accept the trade. Since the agent will trade each of the three last choice nodes, the agent should also trade at each of the second to last nodes.

At the second to last node, which is at the second stage of trades, because they can predict that they will rationally trade at each of the last nodes, they will trade at each of the second to last nodes as well. In the second stage of the upper part of the decision tree, they treat the decision as if they were picking between B-30p, and C-20p. This is because, if, at this part of the decision tree, they choose to trade, they can assume they will also trade at the third stage, and will bring about B-30p. Using the same backwards induction reasoning, they can also reason that if they choose to not trade at this second stage, they will trade in the third stage and end up with outcome C-20p. However as B-30p>C-20p, they will choose to trade at the second stage of the upper part of the decision tree. Similarly, at the second stage bottom part of the decision tree, they predict that trading will lead to C-20p and not trading will lead to A-10p. As C-20p>A-10p, they will choose to trade. At the first choice point, they know that trading will ultimately lead to B-30p and not trading will lead to C-20p. As B-30p>C-20p, they will trade. However, B-30p is worse than their original starting point, B, and an agent acting in accordance to backwards induction will have been successfully money pumped. Rabinowicz has shown that there are numerous other cases in which agents acting in accordance with backwards induction can be money pumped. It is not necessary to outline them here, as the case already outlined has shown that agents acting in accordance with backwards induction are still susceptible to money pumping and so backwards induction is not a suitable solution to money pump problems like the persistent case mentioned above.

However, if you combine my deontic ordering method with a *resoluteness* strategy for money pumps, then you can avoid foreseeable persistent money pumps like the ones Rabinowicz presents. In particular, this strategy involves constructing my deontic ordering between all of the final options in the decision tree, and then *resolutely* selecting a series of decisions that will bring about the outcome for which we have most reason to bring about. That is, we ought to order B-30p, C-20p, A-10p and B using my proposed method and then make the relevant choices to bring about the most choice worthy option. The ordering will depend on the cardinal differences between the options, if cardinal comparisons are possible. The descriptions of two foreseeable money pump cases so far have only involved ordinal axiological comparisons, so I will apply the Copeland method to the ordinal axiological comparisons: A-10p>B, C-20p>A-10p, B-30p>C-20p, B>B-30p and A-10p>B-30p. A-10p 'wins' twice and 'loses' once; B 'wins' once and 'loses' once; C-20p 'wins' once and 'loses' once; B-30p 'wins' once and 'loses' twice. According to my method, you have most reason to select A-10p, and you should trade once in order to achieve it. The axiological orderings are the same in both persistent and non-persistent variants of the money pump cases previously discussed and so my method will select the same option (A-10p) as being most choice worthy option in both variants of the case. In the simple non-persistent version of the money pump case, my method will recommend trading once and foregoing all future trades. In the complex persistent version of the money pump case, it will similarly recommend trading once and rejecting two trades, though it does not matter in which order one trades.

The resoluteness strategy does not avoid money pumping in cases where the agent cannot foresee future options. In such cases, the agent is not able to apply the deontic ordering between to future options as they are not known. Here, the agent can and ought to generate a deontic ordering that between current options and past options when deciding whether or not to trade at each choice node. This can be explained by example. When presented with options B & A-10p, as A-10p>B, she ought to justifiably trade for A-10p. In that case, my deontic ordering claims we have more reason to bring about A-10p than B. After that trade, she is given the option to bring about C-20p instead of A-10p. She can then use the axiological pairwise comparisons between all three options from past and present (B, A-10p and C-20p) to generate a transitive 'more reason to bring about rather than' ordering to determine whether she has more reason to trade or not. Again, which option is most choice worthy depends on how much each option is better than another, if such information is available. If only ordinal comparisons are possible, then the following ordinal axiological relations are relevant to the ordinal ordering: A-10p>B, C-20p>A-10p. C-20p 'wins' once; A-10 'wins' once and 'loses' once, and B 'loses' once. So the method will suggest trading for C-20p. However, once B-30p is offered, and the following axiological comparisons, like before, will be relevant A-10p>B, C-20p>A-10p, B-30p>C-20p, B>B-30p and A-10p>B-30p. A-10p 'wins' twice and 'loses' once; B 'wins' once and 'loses' once; C-20p 'wins' once and 'loses' once; B-30p 'wins' once and 'loses' twice. Of the options now available to the agent, the agent has more reason to select C-20p than B-30p, as C-20p has a higher

choiceworthiness score. This shows that this method will avoid money pumping even when foresight is not possible. While the agent is not able to select the most choice worthy of all options because it is no longer available to them, that could not have been foreseen by stipulation.

This is *similar* to the previously considered commitment principle, but note that it is not susceptible to the same objections as the commitment principle. Firstly, it is able to deontically order the options when all options are presented at once. Secondly, its final deontic ordering of outcomes does not depend on the order in which the options were presented. It therefore is not subject to the *ordering problem*. In the above case, its final deontic ordering still ranks A-10p as most choice worthy, even if it is not able to be selected.

This is not the only case in which it is important to consider not just currently available options but our past and/or future choice contexts in deciding how we ought to act now. This is also evident in cases of pre-commitment, where one eliminates future options (or increases the costs of pursuing those options) in order to achieve one's ultimate ends. For example, problem gamblers sometimes subscribe to casino self-exclusion lists that are intended to prevent them from future gambling and should they engage in gambling, prevent them from obtaining the wins from such gambles. The importance of considering our past and future choice contexts is also reflected in the difference between standard versions of maximizing act consequentialism and rule consequentialism. Rule consequentialism, unlike act consequentialism, allows

for individual acts that do not maximize the good, in order that the overall good is maximized through the following of general rules of behaviour. Similarly, broadening the options that are taken into consideration in my proposed deontic ordering allows one to justifiably forego an option that is better in one's current choice context in order to bring about an overall more choice worthy outcome.

Section V - Objections to my Deontic Ordering Methods

V.A - Violation of Pareto

Broome has noted that my method violates a particular version of the Pareto principle. In particular, it will sometimes claim that we have more reason to bring about X than Y even if Y is better than X in some way and X is not better than Y in any way. Furthermore, sometimes it will claim an option X is the option we have *most* reason to bring about, even though it is pareto worse than another option Y. To see this, consider the following three options, which have cardinal values for each of three axiological criteria: Option 1 (3,3,3), Option 2 (5,0,2) and Option 3 (2,2,2). For comparisons between options 1 and 2 , only the first criterion applies (3 vs. 5), for comparisons between options 2 and 3, only the second criterion applies (0 vs. 2), and for comparisons between options 1 and 3, only the third criterion applies (3 vs. 2). Applying my cardinal choiceworthiness deontic ordering gives option 1 a choiceworthiness score of -1;

option 2 a choiceworthiness score of 0 and option 3 a choiceworthiness score of 1. However option 3 is worse than option 1 by all criteria.

My method violates this pareto condition because in order to generate the deontic ranking my method compares how each option evaluatively compares with every other comparable option and evaluative losses against one option can be outweighed by wins against other options. At first this particular violation of the Pareto principle may seem objectionable. However, it is instructive to remember the type of cases this deontic ordering procedure is responding to - cases where there is cyclicity in > such as when B>A, C>B & A>C. In such cases no matter which option is selected as most choice worthy, there will *always* be a better option instead. As a result, any method that selects an option as most choice worthy will violate the aforementioned pareto principle. It is necessary to select an option as most choice worthy to avoid the moral dilemmas presented by Norcross and Parfit. Given all views that prescribe selecting any option as most choice worthy in such cyclic cases will violate the aforementioned pareto principle, the violation of this version of the pareto principle is not a strong objection to my view.

However, if one believes this is a strong objection to my view, a modified version of my view can avoid the objection[136]. Instead of generating the deontic ordering using all available options, one could remove all options that are pareto inferior to another option in the choice set (i.e. all options that are worse

[136] Thanks to Will Crouch for suggesting this modification on my view.

than another option according to at least one criteria and not better according to any criteria) and then only deontically order the remaining options using my proposed cardinal or ordinal methods. On such a modified version of my view, you have less reason to bring about any of these pareto inferior options than any options that are not pareto inferior to any other options. Applied to Broome's case from the last paragraph, you would have less reason to bring about Option 3 than either Option 1 or 2 because Option 3 is pareto inferior. Option 2 would then be most choice worthy, as it would have a choiceworthiness score of 2, and option 1 would have a choiceworthiness score of -2. As I am not moved by the pareto objection for the reasons I have already outlined, I am not inclined to endorse this modification. I mention this potential modification to show that a view similar to mine can avoid this objection.

V.B - Option Individuation

On my proposed way of constructing a transitive deontic ordering out of a non-transitive betterness ordering, what one ought to do crucially depends on which options one has. However, on different views on the individuation of options, different orderings can arise and this could be objectionable. For example, to return to Temkin's example, consider a case where one of the *L* options could be brought about by two different drugs that are chemically distinct but otherwise identical in all respects (cost, effectiveness, side-effects, etc). If the different drugs that are used are considered different options in the Copeland

method I have suggested, then the number of the pairwise wins and losses will change because there is an additional option and that could change the resulting deontic ordering of the options. It is not 'clone proof' - if addition of evaluatively identical options are taken into account the deontic ordering procedure, the overall deontic ordering can change with this addition. This would be unobjectionable if the additional option of such a drug *is* normatively relevant; however as it is indistinct from one of the other options in every respect but chemical composition, it is implausible that it should affect our deontic ordering in this way. This problem of option individuation can be worsened, when there are many different ways of achieving the same outcome - for example when the outcome considered can be achieved by asking any one of a large number of equally suitable colleagues to do the task.

In order to avoid this problem, there has to be a principled way of individuating the options. One should individuate the options normatively - that is, an option is only considered distinct from another if there is a normative difference between the two options. That will prevent the problems of individuation elucidated above by preventing normatively indistinct options from affecting the deontic ordering. On this view, normatively non-identical but merely similar options *can* make a difference to the deontic ordering. However, I accept this implication - new normatively distinct options should be capable of changing our deontic ordering.

I have assumed throughout this paper that the only difference between the options is axiological. However, different normative theories, particularly non-consequentialist theories, will take normative features other than value as important to how we ought to act. I do not presume here to take a stand on these different theories, but only to show that a transitive deontic ordering can be reached on a principled basis out of a non-transitive axiological ordering. Different normative theories may place further constraints on that ordering - perhaps two acts are axiologically identical, but one is impermissible and the other permissible. Such deontic constraints could be theoretically implemented in a way that is consistent with my view, as my view only states how one can generate a transitive deontic ordering out of a non-transitive axiological ordering, assuming that the differences between all options is only axiological.

It should be noted that my two deontic ordering procedures violate the independence of irrelevant alternatives principle, so that the ordering of options can be changed with the addition of new options. For example, there are cases where we have more reason to bring about A rather than B in choice set {*A*,*B*}, where introducing a third alternative *X*, thus expanding the choice set to {*A*,*B*,*X*}, can cause one to have more reason to bring about B rather than A. This would happen, for example, according to my cardinal choiceworthiness method, if B is better than X by 10 units, A is better than B by one unit, and A & X are axiologically non-comparable. According to the Copeland Method, if X & Y are added to the choice set, expanding it from {*A*,*B*} to {*A*,*B*,*X*,*Y*}, and A>B, B>X, B>Y and all other pairwise axiological relations are non-comparable, then we

have more reason to bring about A than B in choice set {*A*,*B*}; but more reason to bring about B than A in choice set {*A*,*B*,*X*,*Y*}. My two views violate the independence of irrelevant alternatives principle because the deontic ordering of options depends on all of the relevant axiological relations, and so the addition of a new option and its new axiological relations can affect the ordering. In this way, my deontic orderings are context dependent. While many actual agents' preferences violate the independence of irrelevant alternatives principle, the independence of irrelevant alternatives principle has some intuitive plausibility. However it should be noted that the principle's name appears to beg the question against views like mine. According to my view, the addition of new options such as X and Y are *not* irrelevant, as the axiological relations between all options in an agent's choice set are relevant. As I have argued previously, it is a virtue of my view that it takes into account all better than relations. So despite the intuitive cost of violating the principle of 'irrelevant' alternatives, I endorse my two deontic ordering methods if the better than relation is cyclic and non-transitive.

The availability of my two cardinal and ordinal deontic ordering methods shows that justified practical reasoning is possible despite the non-transitivity of the better than relation. It also shows how someone with non-transitive preferences could 'get their preferences in order' (that is, make them transitive) in a justifiable way based on their all of their non-transitive preferences.

V.C - Limitations of Actual Agents

It may be objected that the method I suggest is objectionable because it is more complex than the *knocking out method* - so much so that it is practically impossible to use when comparing certain sets of actual options[137]. This is because it requires many more comparisons than the knocking out method mentioned in Section II.C and sometimes the number of comparisons necessary would be impossible for actual agents with limited and finite cognitive capacities. I concede both of these features of my method, but deny that they are decisive objections. While my method does involve more comparisons, note that just like the simpler knocking out method, my method relies just on pairwise comparisons. Similarly, while my method sometimes requires more comparisons than agents with finite cognitive abilities are able to consider, this is *also* true of the knocking out method. When there are many options, it is impossible for an actual human agent with finite cognitive capacities to compare all of the options and comparisons necessary. Furthermore, while my method requires more comparisons, it should be noted that cases instantiating this particularly troubling complexity are very rare, given they are a subset of the already rare set of cases in which non-transitivity is instantiated.

In addition, it should be stressed that I am not making either the empirical claim that actual agents do reason in this way or the normative claim that they ought

[137] Larry Temkin makes an objection like this to a model that is similar to, but distinct from my deontic ordering methods. See Temkin, *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning,* 468-471.. For more on the important differences between my method and his model, see footnote 130.

to always reason in this way when presented with non-transitive axiological orderings. There may be more efficient ways to reason. Just in the same way that agents do not (and ought not) always use the knocking out method when actually practically reasoning, actual agents ought not always use the two deontic ordering methods I describe. In this paper, I am focusing on the actual underlying fundamental normative principles, not first order practical reasoning prescriptions for actual agents with all of their cognitive, epistemic and motivational limitations[138].

V.D - Are my Two Deontic Ordering Methods Ad Hoc?

It may be objected that using the two deontic ordering methods I have presented to generate a transitive deontic ordering from a non-transitive axiological ordering is ad hoc. In particular, it could be claimed that it is objectionable that a different method of generating a deontic ordering out of an axiological ordering is being used in cases that instantiate non-transitivity of better than in cases that do not instantiate such non-transitivity or cyclicity. Proponents of such an objection claim that there should be one unified method to generate a deontic ordering from all axiological orderings. This objection has some merit. I will argue that we ought to adopt the two methods I have

[138] For those that disambiguate ought into 'objective' and 'subjective' senses, my deontic ordering thus produces the 'objective' ought to do rather than relation rather than the 'subjective' ought to do rather than relation. On this distinction, see Derek Parfit, *On What Matters* Oxford University Press, 2011)..

advocated as a unified method to generate transitive deontic orderings from all evaluative orderings.

The two deontic ordering methods I have advocated can be applied to both transitive and non-transitive axiological orderings to generate a transitive deontic ordering. Furthermore when they are applied to complete axiological orderings that do not instantiate non-transitivity of better than, the methods will generate a rank ordering that is co-extensive with the betterness ordering[139]. To see this, first consider the ordinal Copeland method version, which is used when evaluative cardinal comparisons are not possible between options. Because non-transitivity of better than is not instantiated and the ordering is complete, the axiologically best option (or options) will win the most number of times in pairwise comparison (winning against all worse options), followed by the next axiologically best option (with one less win than the best option when there is only one best option), and so on down to the axiologically worst option (or options) that will lose the most number of times (losing against all better options). To take a particular example where A>B>C>D, A will win 3 times, B will win 2 times and lose once, C will win once and lose 2 times and D will lose 3 times. As such, the rank ordering generated by the Copeland method will be the same as the original axiological ordering. Now consider a complete, transitive axiological ordering where cardinal evaluative comparisons are possible and

[139] However if one uses the cardinal choiceworthiness method, the differences in magnitude between the different options in the deontic ordering may be different from the differences in magnitude between those options in the axiological ordering. The rank order, however, remains the same, so the best, second best, third best and so on correspond with the most choice worthy, second most choice worthy and so on.

apply the cardinal version of my method. The axiologically best option (or options) will not only win the most number of times in pairwise comparison. More importantly for this method, the degree to which that axiologically best option (or options) wins against every worse option will be greater than any other worse option wins against options that are worse than it. This will give the axiologically best option the highest choiceworthiness score on my method. Furthermore, the next best option will get the next best choiceworthiness score and so on until the worst option that gets the lowest score. For example consider A>B>C where A has a value of 10, B has a value of 6 and C has a value of 3. A gets a choiceworthiness score of 11 (4+7); B gets a choiceworthiness score of -1 (-4+3); and C gets a choiceworthiness score of -10 (-7-3). As evident from this example, it should be noted that on the cardinal choiceworthiness method, while the rank ordering is the same, the magnitude of the differences between options is different between the axiological and deontic rankings.

One could use one of my two proposed deontic ordering methods to generate deontic orderings from any set of axiological orderings and the account would thereby be unified. Where the betterness orderings and deontic orderings come apart are in cases that display non-transitivity of better than or certain cases where there is an incomplete axiological ordering. This is because the deontic ordering methods I advocate always generate a complete, transitive deontic ordering. However this is unsurprising, as one of the motivations for my two methods is that they deny that the axiological and deontic orderings are co-extensive. This feature is necessary to enable one to generate a transitive

deontic ordering out of a non-transitive deontic ordering. It is only in these cases of incompleteness or non-transitivity that the rank orderings come apart. It is these very cases that it is implausible that the axiological ordering is co-extensive with the deontic ordering, so it is a virtue of my theory that they come apart according to the non-coextensive thesis.

It ought not to be surprising that deontic reasoning gets somewhat more complex if > is not transitive. I am not arguing that practical reasoning is precisely the same if > is not transitive than if it is transitive. Nor am I arguing that my two proposed deontic ordering methods for generating a transitive deontic ordering from a non-transitive axiological ordering are without theoretical costs. They are more complex than views that suggest better than and more reason to bring about rather than are co-extensive. My two deontic ordering views also violate the independence of irrelevant alternatives principle, a principle that has some intuitive plausibility. However, accepting these costs is far better than being susceptible to the aforementioned money pumps, practical reasoning problems and the Parfit/Norcross moral dilemmas that result from accepting non-transitivity and cyclicity of the better than relation while denying the conditional co-extensiveness thesis. It is far better to adopt the conditional co-extensiveness thesis. Furthermore, if better than is not transitive, we should adopt my two deontic ordering methods to generate complete transitive deontic orderings out of all of our axiological orderings. Even if > is not transitive, it is still possible to reason practically in a justifiable way by

adopting these two deontic orderings, thereby mitigating some of the purported costs of affirming that > is not transitive.

Conclusion

Non-transitivity of the better than relation is supposed to threaten our conception of practical and moral reasoning, making practical dilemmas impossible to avoid, and making agents susceptible to *money pumping*. However I have argued that these threats can be mitigated by accepting the conditional non-coextensive thesis and generating complete transitive deontic orderings out of non-transitive axiological orderings using two different methods I have defended in this paper. I suggested a new deontic principle that prescribes which act we ought to select on the basis of the choice context. All possible pairwise betterness comparisons are made within the available options. If cardinal axiological comparisons are possible (and epistemically available), my cardinal choiceworthiness method uses this cardinal axiological information to order the options in terms of choiceworthiness. If only ordinal axiological comparisons are available, then the Copeland method that does not require this information can be utilized instead. These methods generate a complete, transitive deontic ordering even if the underlying axiological ordering they are based on are either incomplete or non-transitive. By uncoupling the tight link between the axiological and the deontic, some of the supposed costs of non-transitivity of better than can be avoided or mitigated and it is not as

theoretically devastating to accept as some have thought. By denying better than is transitive, while maintaining that more reason to bring about rather than *is* transitive, one can account for important axiological intuitions while maintaining a plausible deontic theory. By adopting my proposed methods for creating a transitive deontic ordering out of a non-transitive evaluative ordering, we can *make betterness behave*.

Bibliography

Almeida, Michael, ed. *Imperceptible Harms and Benefits.* Dordrecht, Holland: Kluwer Academic Publishers, 2000.

Bader, Ralf. "Neutrality and Conditional Goodness." (Unpublished Manuscript).

Kamm, Frances Myra. *Intricate Ethics : Rights, Responsibilities and Permissible Harm.* New York, N.Y. ; Oxford: Oxford University Press, 2007.

Norcross, Alastair. "Contractualism and Aggregation." *Social Theory and Practice* 28, no. 2 (2002): 303-314.

Nozick, Robert. *Anarchy, State, and Utopia.* Basic Books, 1977.

Parfit, Derek. *On what Matters.* Oxford University Press, 2011.

Parfit, Derek. "Justifiability to each Person." *Ratio: An International Journal of Analytic Philosophy* 16, no. 4 (2003): 368-390.

Parfit, Derek. *Reasons and Persons.* Repr. with corrections ed. Oxford: Clarendon Press, 1984.

Rachels, Stuart. "Counterexamples to the Transitivity of Better than." *Australasian Journal of Philosophy* 76, no. 1 (1998): 71-83.

Ross, David. *The Right and the Good.* Oxford University Press, USA, 2003.

Styron, William. *Sophie's Choice.* Random House, 1998.

Temkin, Larry. *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning.* Oxford University Press, 2012.

Temkin, Larry S. "A Continuum Argument for Intransitivity." *Philosophy and Public Affairs* 25, no. 3 (1996): 175-210.