

ARTICLE

Is close enough good enough?

Campbell Brown

London School of Economics, Houghton Street, London WC2A 2AE, UK
Email: c.f.brown@lse.ac.uk

(Received 19 September 2017; revised 13 January 2019; accepted 15 January 2019; first published online 02 May 2019)

Abstract

Should we allow grave harm to befall one individual so as to prevent minor harms befalling sufficiently many other individuals? This is a question of aggregation. Can many small harms ‘add up’, so that, collectively, they morally outweigh a greater harm? The ‘Close Enough View’ supports a moderate position: aggregation is permissible when, and only when, the conflicting harms are sufficiently similar, or ‘close enough’, to each other. This paper surveys a range of formally precise interpretations of this view, and reveals some of the problems they face. It also proposes a novel interpretation which avoids these problems.

Keywords: aggregation; Close Enough View

1. Introduction

You can’t always please everyone. Frequent are the circumstances in which benefiting some people requires burdening others. For example, allowing pesticides may help farmers but hurt beekeepers. How should such conflicts be resolved? One factor that seems surely relevant is the *size* of the harms or benefits involved. We ought to avert the greater harm, or bestow the greater benefit, other things being equal. For instance, if the livelihoods of beekeepers would be destroyed by pesticides, whereas farmers would be only mildly inconvenienced by banning them, then this would speak strongly in favour of a ban. More controversial, however, is another factor: the *number* of harms or benefits. Suppose the farmers are many but the beekeepers few. Does this count against a ban? Could a difference in the number of harms be so great as to outweigh a large disparity in their size? As commonly put, should the numbers count?

Intuitions may be pumped in either direction.¹ Consider first the following well-known example (Scanlon 1998: 235). During a live TV broadcast of the World Cup final, an accident in the transmitter room leaves an engineer trapped under heavy electrical equipment. He is badly injured and suffering painful electrical shocks. Rescuing him would require interrupting the broadcast for 15 minutes.

¹See e.g. Voorhoeve (2014: 64–65).

The match will not be over for another 45 minutes.² Should the rescuers wait till it ends? The trapped engineer, if not rescued immediately, will suffer a harm much greater than the harm to any individual TV viewer of missing 15 minutes of football. On the other hand, the one engineer is vastly outnumbered by the millions of viewers. Many agree with Scanlon that the latter consideration is irrelevant. Regardless of how many are watching the broadcast, we ought to save the engineer straightaway. The claim, to be clear, is not merely that the number of viewers required to justify delaying the rescue is very, very large, nor even that it is so large that, given contingent facts about the population of our planet and so on, it could never in practice be reached. Rather, it is the stronger claim that there is *no such number*, not even a practically unattainable one. For *any* number *n*, we should not allow one person to suffer in agony for 45 minutes in order to allow *n* people to each enjoy 15 minutes of minor pleasure, were it possible to do this.

We should, of course, be cautious in trusting our intuitions, especially in cases like this. The number of World Cup viewers may already be too large for us to comprehend. But Scanlon's claim applies to even larger, unimaginably large numbers. As Broome (2004: 57–59) points out, intuitions about such large numbers are not generally reliable.³ Still, it would be rash simply to dismiss these intuitions out of hand. It is worthwhile to consider whether they can be accommodated by an otherwise plausible view.

Consider now a variation of Scanlon's example. In this case, there is no accident but instead a fault prevents the broadcast from reaching some viewers. Repairing the fault would require interrupting the broadcast to *all* viewers for 15 minutes. As before, the match still has 45 minutes to run. Should the fault be fixed? Here it seems the numbers are relevant. Were we to learn the fault is affecting only a handful of viewers, then it would seem reasonable to leave repairs till after the match. Though the harm to the affected viewers would be greater (missing 45 minutes of football, as opposed to only 15), these viewers are but a tiny proportion of the TV audience. On the other hand, if the vast majority of viewers were affected, it would be right to fix the fault immediately. In intermediate cases, it may be less clear what is the right thing to do. But the numbers involved would be a relevant factor.

Examples like these may lead us to seek a middle ground, a 'moderate' view according to which numbers count in some cases but not others. But what makes this difference? The following seems a natural answer. Each of these cases involves a conflict between two harms. These 'conflict' in the sense that it is impossible to prevent both. In the second case, where numbers *do* seem to count, the conflicting harms are very similar. There is not a great difference between being deprived of football viewing for 45 minutes versus 15 minutes. In the first case, on the other hand, where numbers *do not* seem to count, the harms are very dissimilar. There is quite a significant difference between being deprived of football viewing versus suffering intense physical pain.⁴

²I have slightly adjusted the timing from Scanlon's original example, to avoid a minor technical hitch.

³See Voorhoeve (2014: 75–76) for a counter-argument to Broome.

⁴Some may object that being deprived of pleasure is no harm at all; the absence of a benefit is not the same as the presence of a harm. Personally, I am disinclined to put much weight on this distinction. It rests on the selection of a 'baseline', relative to which benefits and harms are calculated. But this is entirely a contextualist

Scanlon suggests a view along these lines. He describes it as follows.

If one harm, though not as serious as another, is nonetheless serious enough to be morally 'relevant' to it, then it is appropriate, in deciding whether to prevent more serious harms at the cost of not being able to prevent a greater number of less serious ones, to take into account the number of harms involved on each side. But if one harm is not only less serious than, but not even 'relevant to', some greater one, then we do not need to take the number of people who would suffer these two harms into account in deciding which to prevent, but should always prevent the more serious harm. (Scanlon 1998: 239–240)

Parfit calls this the 'Close Enough View'. He defines it as follows.

When burdens to different people are close enough in size, one greater burden could be morally out-weighted by a sufficient number of lesser burdens. (Parfit 2003: 378)

As Parfit notes, however, this does not yet ensure moderation. Even Utilitarians, for example, may accept the Close Enough View as defined, because they believe that *all* burdens are close enough. To capture the moderate spirit of the Close Enough View, we need, therefore, to add that some but not all burdens are close enough.

Neither Scanlon nor Parfit develops this view in precise detail. Voorhoeve (2014) proposes a more detailed version, which he calls 'Aggregate Relevant Claims'. But he explicitly restricts this view to a limited range of cases (Voorhoeve 2014: 67). My primary aim in this paper is to remedy these deficits, to investigate the prospects for developing a formally precise and comprehensive version of the Close Enough View.⁵ This is an important project. A view that is appealing when applied to simple cases may turn out to have unpalatable implications when extended further.

I proceed as follows. I begin by defining a formal framework with which to represent distributive views. These are formally represented by 'choice functions'. I then delineate a subclass of choice functions, which I call 'moderate'. The candidate interpretations of the Close Enough View to be evaluated are drawn from this subclass. I then evaluate some specific candidates. These I divide into two camps: 'individualist' and 'collectivist'. An individualist view compares harms only between *individuals*, whereas a 'collectivist' view compares harms also between *collections* of individuals. Each of the candidates considered is a generalization of a familiar 'extreme' view (one that requires aggregation either always or never).

I argue that some of the candidates are unacceptable, because they have problematic implications. In particular, I reject some choice functions which may be seen as natural extensions of Voorhoeve's proposal. But I also propose another candidate, the

matter: what may properly be called a harm in one context may no less properly be called the withholding of a benefit in another. There is no further fact as to which context is correct. But, in any case, this is an inessential feature of the examples. We could devise variants where it is more clearly harm on both sides, and intuitions, I conjecture, would be much the same.

⁵My discussion is still limited in two important respects: I do not consider uncertainty or variable population cases. More on this below.

combination of conditions I call ‘Individualist-Collectivist Hybrid’ and ‘Transitive Closure Maximization’, and show that this avoids the problems of the others. This family of choice functions, I argue, provides the most promising interpretation of the Close Enough View. However, as I also point out, Individualist-Collectivist Hybrid has features that some may find unappealing. So some may instead treat this as an argument for extremism. In any case, my investigation will reveal some of the costs of moderation. Whether these costs are worth paying I will not try to definitively answer.

2. Preliminaries

My aim in this section is to delineate a particular subclass of distributive views – or choice functions, as they will be formally represented – from which the candidates to be evaluated in the remainder of the paper will be drawn. Given the vastness of the space of possible choice functions, it is helpful to narrow our focus in this way. To do this, I need to define a general formal framework. This will involve formally defining the concept of a ‘harm’ and of what I call its ‘significance’. The relation of being ‘close enough’ is then defined in terms of significance. Finally, I show that, within the class of choice functions defined, the degree to which a choice function is friendly to aggregation is represented by a single parameter.

2.1. Basic framework

I begin by laying out some assumptions underlying the formal framework to be adopted. I assume that the permissibility of an action (or policy, etc.) depends solely on the value of its outcome in comparison with the those of its alternatives. I set aside uncertainty, so that each available action is associated with a single outcome. In effect, the options between which one chooses are outcomes. I assume, further, that the value of an outcome supervenes on the distribution of wellbeing it contains: if two outcomes contain the same distribution, then they must be equally good. In short, I assume a ‘consequentialist’ and ‘welfarist’ framework.

By the ‘population’ of an outcome, I mean the set of individuals who exist in that outcome. I need to allow outcomes with arbitrarily large populations (though, to avoid unnecessary complications, I assume populations must be finite). The Close Enough View claims, for example, that preventing one death is better than curing ‘any number’ of mild headaches. To represent this claim in my framework, it must therefore be possible for any number of people to exist. However, again for simplicity, I consider only ‘same people’ choices, where all available outcomes share a common population.

Let $\mathbb{N} = \{1, 2, \dots\}$ be the set of all possible individuals.⁶ An *outcome* is a function $a: \mathbb{N} \rightarrow \mathbb{R}_+ \cup \{\Omega\}$, such that the set $\mathbb{N}_a = \{i \in \mathbb{N} : a(i) \neq \Omega\}$, called the *population* of a , is finite. Thus $a(i)$ is either a non-negative real number, representing the wellbeing of individual i , or Ω , representing non-existence.⁷ I assume wellbeing is measurable

⁶For simplicity, I assume these are countably infinite.

⁷By \mathbb{R}_+ I mean the non-negative reals. The reason for excluding negative wellbeing is that the ‘prioritarian’ views I shall consider do not satisfy ‘ratio-scale invariance’ when negative wellbeing is included, as shown e.g. by Brown (2007).

and interpersonally comparable on a ratio scale. For neatness, I shall sometimes write a_i instead of $a(i)$.

A *choice situation* is a non-empty, finite set of outcomes A with a common population (i.e. $N_a = N_b$ for all $a, b \in A$). Let \mathcal{A} be the set of all choice situations. The purpose of a distributive view is to determine, for a given choice situation, which of the options may permissibly be chosen. Such a view I represent by a *choice function* $\gamma : \mathcal{A} \rightarrow \mathcal{A} \cup \{\emptyset\}$ such that $\gamma(A) \subseteq A$. The options in $\gamma(A)$ are those which may permissibly be chosen from A . Thus a choice function may contract the agent's choice situation by eliminating some options as impermissible. (I have here left open the possibility that $\gamma(A)$ is empty, so that all options are eliminated. I return to this issue below.)

It will be useful to impose a further condition on choice functions. Let σ be a permutation of \mathbb{N} . Then I shall write σa to denote the composition of a with σ (i.e. $\sigma a = a \circ \sigma$). So $\sigma a(i) = a(\sigma(i))$ for all $i \in \mathbb{N}$. And I shall write σA to denote the choice situation $\{\sigma a : a \in A\}$. The further condition may now be stated as follows: $\gamma(\sigma A) = \sigma \gamma(A)$. This requires that the choice function is, in a certain sense, 'anonymous'. Suppose we begin with one choice situation and then transform this by uniformly replacing some of the individuals. If individual i is replaced by individual j in one outcome, then i must also be replaced by j in every other outcome. The distribution of wellbeing levels in each outcome remains the same, except that which particular individuals have these levels may change. If j replaces i , then j must receive i 's wellbeing level in every outcome. This condition requires that which outcomes are permissible should remain unchanged by this transformation.

This seems incredibly plausible. Moreover, it also has the expository benefit of simplifying some of the definitions to be given below. I shall write (x_1, x_2, \dots, x_n) to denote the outcome a such that

$$a_i = \begin{cases} x_i & \text{if } i \leq n \\ \Omega & \text{if } i > n \end{cases}$$

Further, where n consecutive places in such a sequence have the same value x , I shall abbreviate this as $n \times x$. Thus, for example, $(3 \times x) = (x, x, x)$ and $(x, 2 \times y) = (x, y, y)$. The set of choice situations involving outcomes that can be denoted in this way is a proper subset of all choice situations. However, given the anonymity condition, it follows that a choice function must be reducible to the choice situations in this subset. If the choice function satisfies a condition that is explicitly defined only for situations in this subset, then it must implicitly satisfy this condition for all situations. This fact may be exploited to give simpler definitions of such conditions.

A binary choice situation is one containing only two options. In this context, a choice function may be represented by a binary relation. For a choice function γ , let \geq_γ be a binary relation such that $a \geq_\gamma b$ if and only if $a \in \gamma(\{a, b\})$. Thus $a \geq_\gamma b$ may be read as saying that, according to γ , it is permissible to choose a rather than b (when these are the only options). Let $>_\gamma$ and \sim_γ be, respectively, the asymmetric and the symmetric parts of \geq_γ . So we have $a >_\gamma b$ if and only if $\gamma(\{a, b\}) = \{a\}$, and $a \sim_\gamma b$ if and only if $\gamma(\{a, b\}) = \{a, b\}$. Thus $a >_\gamma b$ may be read as 'it is obligatory to choose a rather than b ', and $a \sim_\gamma b$ as 'it is optional to choose either a or b ' (again

when these are the only options). Occasionally, I shall speak a little loosely, expressing $a \succ_{\gamma} b$ as ‘ a is better than b ’. But I should stress that this is not the official definition. For one thing, it is plausible to think that ‘better than’ necessarily expresses a transitive relation. But it is less clear that \succ_{γ} , as defined above, must be transitive (more on this below, Section 4.1). In what follows, I suppress the subscript γ , writing simply \succeq , but it should be borne in mind that this relation is defined relative to a choice function.

2.2. Harms and their significance

The Close Enough View, recall, is a proposal for adjudicating between conflicting harms. So I need to explain how harms are represented in the formal model. I assume a counterfactual definition of harm: a person is harmed to the extent that her wellbeing is less than it would have been otherwise. This assumes, of course, that there is some determinate fact as to what would have happened otherwise. Happily, for binary choices, this is no problem. The options in a choice situation are assumed to be exhaustive: one of them must be chosen. Thus, if there are only two options, then whichever option is chosen, the other is what would have happened otherwise. (For obvious reasons, this is less straightforward in non-binary choices. I return to this below.) Suppose the two options are a and b , and a is the chosen option. Then, as a result, a_i is the wellbeing of individual i , and b_i the wellbeing i would have had otherwise. So i is harmed in a if and only if $a_i < b_i$, and the size of the harm is $\max\{0, b_i - a_i\}$. (Notice, I do not allow ‘negative’ harms, and a harm of size zero is really no harm at all.)

I distinguish between the *size* of a harm and its *significance*. By ‘significance’ I mean roughly the strength of moral reason we have to prevent the harm. On some views, significance simply reduces to size: we ought always to prevent the greater harm. But others reject this. Consider, for example, the following pair of outcomes.

	1	2
a	9	100
b	10	98

Here individual 1 is harmed in a , and individual 2 is harmed in b . Although the latter harm is greater in size ($10 - 9 < 100 - 98$), we may nonetheless believe this harm is less significant, because the other, smaller harm falls on a worse-off individual, who will have less wellbeing in any case. On this view, some priority should be given to preventing harms to the worse off. I want to leave open the possibility of such ‘prioritarian’ views.

The significance given to a harm by a distributive view, represented here by a choice function γ , is revealed by the choices γ prescribes. For the following definitions, let a and b be outcomes with a common population, which includes at least two individuals, i and j , such that $a_i < b_i$ and $b_j < a_j$. So i is harmed in a , and j in b . Then I shall say that the harm to i in a is *at least as significant* as the harm to j in b , according to γ , if and only if

$$(b_i, b_j) \succeq (a_i, a_j)$$

Thus, to compare the relative significance of two harms, we consider a pair of options in which only these harms are present. Whichever harm we ought to prevent in this situation is the more significant.⁸

In some cases, a distributive view may say that one harm is ‘infinitely more’ significant than another, in the sense that preventing one individual’s suffering this harm is better than preventing any number of individuals’ suffering the other harm. In this case, I say that the more significant harm ‘trumps’ the less significant. That is, the harm to i in a *trumps* the harm to j in b , according to γ , if and only if, for any $n > 0$,

$$(b_i, n \times b_j) \succ (a_i, n \times a_j)$$

Finally, I say two harms are ‘close enough’ if and only if neither harm trumps the other. That is, the harm to i in a is *close enough* to the harm to j in b , according to γ , if and only if, for some $n, m > 0$,

$$(b_i, n \times b_j) \not\succeq (a_i, n \times a_j) \text{ and } (m \times a_i, a_j) \not\succeq (m \times b_i, b_j)$$

Notice, by definition, ‘close enough’ is a symmetric relation. This may seem a departure from ordinary usage. For example, we might say that a replica of a guitar is close enough to the original, but it would be odd to say that the original is close enough to the replica. In response, I should stress that ‘close enough’ is here used as a semi-technical term.⁹

2.3. Regularity

It will be helpful to narrow my focus to a subclass of choice functions. These are choice functions according to which the relations of significance have a certain simple structure, defined by the following two conditions.

First, the comparative significance of harms is determined by a strictly increasing, and concave function f .¹⁰ So significance is determined not by the difference in wellbeing levels, but rather by the difference in f -weighted wellbeing levels. That is, the harm to i in a is at least as significant as the harm to j in b if and only if

$$f(b_i) - f(a_i) \geq f(a_j) - f(b_j)$$

⁸Some may object to this reduction on the grounds that the significance of the harm to one individual may partly depend on the harms to other individuals. In response to this objection, I would stress that ‘significance’ is being used here as a technical term, which need not perfectly match its ordinary meaning. I could perhaps say instead something like ‘intrinsic significance’ to allay such worries. But that would be too messy.

⁹And in any case, the example might be explained by conversational implicature. It might be true that the original is close enough to the replica, but nonetheless infelicitous to say so.

¹⁰By saying that f is strictly increasing I mean that it slopes upwards, i.e. if $x > y$ then $f(x) > f(y)$. By saying that f is concave I mean that it does not curve upwards; either it is linear or it curves downwards. In the latter case, I say f is *strictly* concave. More precisely, f is concave if, for any x, y and any $\lambda \in [0, 1]$,

$$f(\lambda x + (1 - \lambda)y) \geq f(\lambda x) + f((1 - \lambda)y).$$

For the definition of strict concavity, replace \geq with $>$ in the above.

This allows, but does not require, some degree of priority to the worse off. If f is linear, then significance reduces to size. But if f is strictly concave, then a smaller harm may be more significant than a larger harm. This condition does, however, exclude absolute priority to the worse off. Significance must at least be influenced by size, even if not exclusively determined by it. This seems reasonable. It rules out distributive views according to which, for example, a worse off person's suffering a brief mild headache is more significant than a better off person's dying. But such extreme views hold little appeal from a moderate perspective.

Second, whether one harm trumps another depends on the *ratio* of their 'significances'. More specifically, there exists a ratio $\rho \in [0, 1]$ such that the harm to i in a trumps the harm to j in b if and only if

$$\frac{f(a_i) - f(b_j)}{f(b_i) - f(a_i)} < \rho$$

For example, suppose f is linear and $\rho = 1/2$. Then, for a smaller harm to be close enough to a larger harm, the smaller must be at least half as great as the larger. If it is only, say, one third as great, then it is trumped by the larger harm.¹¹

A choice function that satisfies these two conditions I call, for lack of a better term, 'regular'. That is, a regular choice function satisfies the following condition.

Regularity. There exists a strictly increasing and concave function f and a ratio $\rho \in [0, 1]$, such that, for any $x, y, x^+, y^+ \in \mathbb{R}_+$, with $x < x^+, y < y^+$,

1. $(x^+, y) \geq (x, y^+)$ if and only if $f(x^+) - f(x) \geq f(y^+) - f(y)$;
2. there exists an $n > 0$ such that $(n \times x, y^+) \not\geq (n \times x^+, y)$ if and only if

$$\frac{f(x^+) - f(x)}{f(y^+) - f(y)} \geq \rho$$

I believe Regularity is a plausible condition. As we will see, it is compatible with a broad range of views. In a fuller treatment, I could offer an argument for this condition. But that would be beyond the scope of this paper. By restricting the candidates in this way, I do risk overlooking an irregular choice function that provides a more plausible interpretation of the Close Enough View. Nonetheless, I believe the significant gain in simplicity warrants this risk. Hereafter, by 'choice function' I refer only to regular choice functions.

2.4. Moderation

Choice functions may be classified along a spectrum of attitudes to aggregation. A moderate view, as I said earlier, is one that allows aggregation sometimes but not always. In the framework defined above, this means that more significant harms trump less significant harms in some cases but not in others. That is, a moderate choice function satisfies both of the following conditions.

¹¹The question of how this ratio is determined I will not address here. On this issue I recommend Voorhoeve (2014).

Moderate Aggregation. In some cases, a more significant harm *does not* trump a less significant harm.

Moderate Non-Aggregation. In some cases, a more significant harm *does* trump a less significant harm.

By assuming Regularity, I have not ruled out extremism. For example, if $\rho = 1$, then Moderate Aggregation is violated. This would be a form of extreme ‘non-aggregationism’. On the other hand, if $\rho = 0$, then Moderate Non-Aggregation is violated. This would be a form of extreme ‘aggregationism’. Between these two extremes, we have moderate views according to which $0 < \rho < 1$. Thus we may treat ρ as a parameter measuring the extent to which a choice function is non-aggregationist. The greater is ρ , the more non-aggregationist is the choice function (at least when holding f constant).

Selecting a particular f and ρ determines the choice function only for a limited class of choice situations. This includes only binary choices between outcomes in which harms are uniform (i.e. all harmed individuals in an outcome suffer the same harm). The question, then, is how to extend the choice function beyond this limited class. I take this in two steps. First, I consider binary choices with non-uniform harms. Then I consider non-binary choices.

3. Binary choices

I turn now to evaluating some candidate interpretations of the Close Enough View, in the context of binary choice only. Before discussing specific candidates, however, I need to address an objection that applies to all such views. Indeed, this objection applies to all views that accept Moderate Non-Aggregation, including also extreme non-aggregationist views. Having dealt with this objection, I then consider some collectivist candidates, before turning to some individualist candidates.

3.1. Anonymity

The objection, in brief, is that Moderate Non-Aggregation is jointly inconsistent with two further conditions, both of which, initially at least, may seem very plausible. As I argue, however, one of these conditions is not so plausible after all, and supporters of Moderate Non-Aggregation may feel comfortable rejecting it.

The first condition is commonly called ‘Anonymity’. The basic idea is that the identities of individuals should not matter. All that should matter is the ‘anonymized’ distribution of wellbeing. Suppose one outcome can be obtained from another merely by rearranging the ‘name-tags’ attached to wellbeing levels. Then the choice between these outcomes must be a matter of moral indifference. I shall give this condition another name, ‘Permutation Invariance’, because, for reasons given below, I am sceptical about whether it is merely a requirement of anonymity. It may be defined as follows.

Permutation Invariance. $a \sim \sigma a$.¹²

The second condition I call ‘Non-Arbitrariness’. This requires that if the most significant harm in a is trumped by the least significant harm in b , then we ought to choose a . This seems very plausible. Imagine, for example, a view that implies each of the following claims.

1. We ought to prevent one death rather than prevent any number of broken arms.
2. We ought to prevent one death rather than prevent any number of headaches.
3. For sufficiently large numbers n and m , we ought to prevent n broken arms and m headaches rather than prevent one death.

Such a view seems quite arbitrary. If aggregation is not appropriate when death conflicts either with broken arms or headaches alone, how can it be appropriate when death conflicts with some combination of the two? This condition rules out such views. Formally, it is defined as follows. Let H_{ab} be the set of all individuals harmed in a , relative to b (i.e. those who will be harmed if b is chosen rather than a). So $H_{ab} = \{i: a_i < b_i\}$. Then we have:

Non-Arbitrariness. $b \not> a$ only if

$$\max_{i \in H_{ba}} f(a_i) - f(b_i) \geq \rho \min_{i \in H_{ab}} f(b_i) - f(a_i)$$

That Moderate Non-Aggregation, Permutation Invariance and Non-Arbitrariness are jointly inconsistent can be shown as follows. Suppose $\rho > 0$, as required by Moderate Non-Aggregation. Then there must exist a positive integer m such that

$$f(1/m) - f(0) < \rho(f(1) - f(0))$$

Without loss of generality, suppose $m = 2$. (It should be clear that the proof generalizes to any positive integer $m > 1$.) We may then consider the following outcomes.¹³

	1	2	3
a	0	1/2	1
b	1/2	1	0

Clearly Permutation Invariance implies $a \sim b$.¹⁴ But it follows from the definition of m , and the concavity of f , that the one harm in b (to individual 3) trumps both of the harms in a (to individuals 1 and 2).¹⁵ So Non-Arbitrariness then implies $a > b$.

¹²This condition must be restricted to permutations that are, as we might say, \mathbb{N}_a -preserving. That is, we must have $i \in \mathbb{N}_a$ if and only if $\sigma(i) \in \mathbb{N}_a$. This ensures that $\{a, \sigma a\}$ is a choice situation, i.e. with a fixed population.

¹³Parfit (2003: 383, n. 16) discusses a case like this, and Voorhoeve (2014 82–84) discusses a variant of Parfit’s case.

¹⁴Let σ be the permutation of \mathbb{N} such that $\sigma(1) = 2, \sigma(2) = 3, \sigma(3) = 1$, and $\sigma(i) = i$ for $i > 3$. Then $b = \sigma a$.

¹⁵From the definition of m we have $f(1/2) - f(0) < \rho(f(1) - f(0))$. And since f is concave, $f(1) - f(1/2) \leq f(1/2) - f(0)$. So we also have $f(1) - f(1/2) < \rho(f(1) - f(0))$.

In my view, this is not a compelling objection to Moderate Non-Aggregation. Moderates may instead reject Permutation Invariance. This condition is intended to capture the ideal of ‘moral equality’: all individuals’ interests should be given equal weight, and no one should be favoured simply because of who she is. But Permutation Invariance actually requires more than this. Consider the following two pairs of outcomes.

	1	2	3
<i>a</i>	0	1/2	1
<i>b</i>	1/2	1	0

	1	2	3
<i>c</i>	1	1/2	0
<i>d</i>	0	1	1/2

A non-aggregationist view that judges *a* better than *b* will also judge *c* better than *d*. But this shows that there is no violation here of moral equality. This view favours preventing harm to individual 3 in the first choice (*a* vs *b*), not because of who she is, but because of her situation. Moving to the second choice (*c* vs *d*), individual 1 is now in the same situation that individual 3 was in, and now this view favours preventing harm to individual 1. Although this view violates Permutation Invariance, it cannot fairly be accused of weighing the interests of some individuals more highly than those of others.

The ideal of moral equality is better captured by a weaker condition. Let a *transposition* be a permutation τ such that $\tau \circ \tau$ is the identity function, i.e. $\tau(\tau(i)) = i$ for all *i*. A transposition merely ‘swaps’ the positions of one or more pairs of individuals, so that applying the permutation twice swaps them back again, returning us to where we started. Then the condition I propose is the following.

Transposition Invariance. $a \sim \tau a$.

Suppose we are choosing between outcomes *a* and τa . Then, for any individual *i* who suffers a harm in *a*, there is another individual, $\tau(i)$, who suffers precisely the same harm in τa . The harms in *a* are identical to the harms in τa , except for the identities of who is harmed. Therefore, if all individuals’ interests are weighed equally, we must be indifferent between *a* and τa . In this way, Transposition Invariance captures moral equality.¹⁶

Clearly, Permutation Invariance implies Transposition Invariance, but not vice versa (because the transpositions are a proper subset of the permutations). The permutations involved in the examples above are not transpositions, and

¹⁶It should be noted that every permutation may be decomposed into a sequence of transpositions. So, if \sim is transitive, then Transposition Invariance collapses into Permutation Invariance. However, as I point out below, non-aggregationists are already committed to denying the transitivity of \sim .

therefore Transposition Invariance does not apply in these cases. Moderate Non-Aggregation is consistent with the conjunction of Transposition Invariance and Non-Arbitrariness.¹⁷ So we may accept a plausible form of anonymity while rejecting extreme aggregation.

3.2. Collectivism

Having dealt with that general objection, I may now move on to evaluating some specific candidates. The first I consider is a form of collectivism. I first define an extreme form of collectivism, and then show how this may be generalized to accommodate moderation.

This extreme form of collectivism tells us to minimize the total significance of unprevented harms; or, in other words, to choose the outcome in which the harms have the least total significance. This class of choice functions may be defined as follows.

Extreme Collectivism. $a \geq b$ if and only if

$$\sum_{i \in H_{ab}} f(b_i) - f(a_i) \leq \sum_{i \in H_{ba}} f(a_i) - f(b_i)$$

This is in fact equivalent to a familiar class of choice functions, sometimes called ‘Generalized Utilitarianism’. If f is linear, it is equivalent to the view that we should maximize total wellbeing (a.k.a. Utilitarianism).¹⁸ If f is strictly concave, it is equivalent to the view that we ought to maximize total priority-weighted wellbeing (a.k.a. Prioritarianism).

Extreme Collectivism implies $\rho = 0$, and so is incompatible with Moderate Non-Aggregation. How might Extreme Collectivism be generalized so as to allow moderation? The generalization I shall consider is inspired by Voorhoeve’s ‘Aggregate Relevant Claims’ (Voorhoeve 2014: 66). As the name suggests, Voorhoeve frames his proposal in terms of *claims*, rather than harms. But this difference seems merely terminological. For Voorhoeve, when a person is harmed in an outcome, she has a *claim* against this outcome (or, as he might prefer to say, a claim in favour of the other outcome). The strength of a claim depends both on the size of the harm and on how badly off is the claimant. Thus, where I talk of preventing more significant harms, Voorhoeve talks instead of satisfying stronger claims. So far as I can tell, this amounts to essentially the same thing. I should also note that Voorhoeve explicitly restricts his proposal to a limited class of choice situations, which does not include some of those I discuss below. The view I suggest seems the most natural extension

¹⁷Given the definition of a choice function, which includes an anonymity condition, Transposition Invariance is equivalent to the following: for any \mathbb{N}_a -preserving transposition σ , $\gamma(\{a, \sigma a\}) \neq \emptyset$.

¹⁸Because $H_{ab} \cap H_{ba} = \emptyset$, and because $a_i = b_i$ for all $i \notin H_{ab} \cup H_{ba}$, we have

$$\begin{aligned} \sum_{i \in H_{ab}} f(b_i) - f(a_i) - \sum_{i \in H_{ba}} f(a_i) - f(b_i) &= \sum_{i \in H_{ab}} f(b_i) - f(a_i) + \sum_{i \in H_{ba}} f(b_i) - f(a_i) = \sum_{i \in I} f(b_i) - f(a_i) \\ &= \sum_{i \in I} f(b_i) - \sum_{i \in I} f(a_i) \end{aligned}$$

Thus $\sum_{i \in H_{ab}} f(b_i) - f(a_i) \leq \sum_{i \in H_{ba}} f(a_i) - f(b_i)$ if and only if $\sum_{i \in I} f(a_i) \geq \sum_{i \in I} f(b_i)$.

of his proposal to these other cases. Nonetheless, it should be noted that the problems I raise are not strictly speaking problems for Aggregate Relevant Claims.

The basic idea is that, rather than aggregating all harms, we aggregate only ‘relevant’ harms. A harm is relevant if and only if it is not trumped by any conflicting harm. Irrelevant harms, says Voorhoeve, should ‘play no role’ in determining which outcome to choose (Voorhoeve 2014: 66). Thus his proposal is like Extreme Collectivism, except it disregards all irrelevant harms. That is, it tells us to minimize the total significance of all *relevant* unprevented harms. Let R_{ab} be the set of individuals who are *relevantly* harmed in a relative to b .¹⁹ So we have

$$R_{ab} = \{i \in H_{ab} : \forall j \in H_{ba} f(b_i) - f(a_i) \geq \rho(f(a_j) - f(b_j))\}$$

Then we may define the following condition.

Generalized Collectivism. $a \geq b$ if and only if

$$\sum_{i \in R_{ab}} f(b_i) - f(a_i) \leq \sum_{i \in R_{ba}} f(a_i) - f(b_i)$$

Generalized Collectivism is a generalization of Extreme Collectivism, in the sense that the choice functions satisfying the former are a proper superset of those satisfying the latter. If $\rho = 0$ then all harms are relevant, $R_{ab} = H_{ab}$, so in this case Generalized Collectivism is equivalent to Extreme Collectivism. But Generalized Collectivism is also compatible with the conjunction of Moderate Aggregation and Moderate Non-Aggregation (i.e. we may have $0 < \rho < 1$).

Does Generalized Collectivism provide a satisfactory interpretation of the Close Enough View? I shall argue that it does not. Consider the following condition.

Permutation Pareto. If a is Pareto-superior to σb then $a > b$.²⁰

Friends of Moderate Non-Aggregation may reject Permutation Pareto for the same reason that I argued above they may reject Permutation Invariance. However, suppose we weaken Permutation Pareto in the same way that suggested we should weaken Permutation Invariance. This gives us the following condition.

Transposition Pareto. If a is Pareto-superior to τb then $a > b$.

Transposition Pareto and Generalized Collectivism together imply Extreme Collectivism. That is, given Transposition Pareto, our generalization collapses into its extreme form. Suppose as before, without loss of generality, that

$$f(1/2) - f(0) < \rho(f(1) - f(0))$$

Then consider the following outcomes.

¹⁹Strictly speaking, this should have further subscripts, $R_{abf \rho}$, indicating that this set is defined also relative to a function f and ratio ρ . But I omit these for neatness.

²⁰ a is Pareto-superior to b if and only if $\mathbb{N}_a = \mathbb{N}_b$, $a_i \geq b_i$ for all $i \in \mathbb{N}_a$, and $a_i > b_i$ for some $i \in \mathbb{N}_a$.

	1	2	3
<i>a</i>	0	1	1/2
<i>b</i>	1	0	0

Since the harm to individual 3 in *b* is deemed irrelevant, Generalized Collectivism implies $a \sim b$. But clearly Transposition Pareto implies $a > b$, since an outcome that is Pareto-inferior to *a* may be obtained from *b* merely by swapping wellbeing levels between individuals 1 and 2.

This seems a troubling result for moderate forms of Generalized Collectivism. What reason can be given for allowing the harm to individual 3? One might say that allowing this harm is necessary in order to prevent the much more significant harm to individual 1. But in allowing the harm to individual 3, we also allow a harm to individual 2 which is equally as significant as the harm that we prevent. So far as ‘relevant’ harms are concerned, there is a perfect tie. We should then be able to use ‘irrelevant’ harms to break this tie. But Generalized Collectivism denies these harms even this minor tie-breaking role. Or, to put the point differently, these harms should not be treated as *entirely* irrelevant, but rather should be given the minimal relevance of tie-breakers.

Someone might object to this argument with an example like the following. Suppose we can save the life of only one of two people, Anne or Bob, where the harm of dying is precisely the same for both. So we have a tie. Before tossing our coin, however, we realize that the cost of saving Anne’s life is slightly more. If we choose to save Bob instead, then we will save £1, which we may use to buy an ice-cream for Cathy. So if we save Anne’s life, then, as well as Bob, Cathy will also be harmed: she will miss out on the ice-cream. However, it may seem absurd to use such a trivial harm to break a tie involving so serious a harm as death.²¹

I have two responses to this objection. First, it points to the issue of vagueness in interpersonal comparisons of wellbeing. The formal framework adopted here assumes that we can make these comparisons in a very precise way. In practice, however, it is very unlikely that we could ever be confident that the harm of death would be precisely the same for Anne and Bob. Rather, we could say only that these harms are equal within a certain margin of error. Since the small harm to Cathy is within this margin of error, adding it to the scales makes no practical difference. In practice, such minor harms may be irrelevant. On the other hand, if we can imagine occupying a ‘god-like’ perspective from which we can determine that the harms to Anne and Bob really are precisely the same, then it seems to me that it would be wrong, from this perspective, simply to disregard Cathy.

Second, it is important to recognize that wrongness comes in degrees. Some wrong actions may be less wrong than others. Now consider a different example which involves only Cathy. We can either give her the ice-cream or destroy it. It would be a gratuitous harm, and therefore wrong, to destroy the ice-cream. But it would be only a minor wrong. Likewise, in the original case, it may be only a comparably minor wrong to save Anne’s life rather than Bob’s.

So Generalized Collectivism, I conclude, is unacceptable. But it may be modified to avoid this problem. Whereas Extreme Collectivism gives irrelevant harms too

²¹Kamm (2005: 13–14) discusses similar examples.

much weight, Generalized Collectivism gives them too little. A solution, then, may be a combination of these two conditions in which the latter has lexical priority over the former. This gives us the following condition.

Lexical Generalized Collectivism. First, minimize the total significance of *relevant* unprevented harms (as in Generalized Collectivism). Second, in the case of a tie, minimize the total significance of *all* unprevented harms (as in Extreme Collectivism).

Lexical Generalized Collectivism is compatible with the conjunction of Transposition Pareto, Moderate Aggregation and Moderate Non-Aggregation.²² A general problem with lexical priority, however, is that it allows very small differences to have very great weight. This can be seen in the case of Lexical Generalized Collectivism. If there is no tie to break, then irrelevant harms are again entirely ignored. This may seem implausible in cases where the more relevant harms, though not *strictly* tied, are very close to being so.

Consider a variation of the World Cup example. In this case there are two trapped engineers, but we can rescue only one at a time. One of the engineers, as before, can only be rescued by interrupting the broadcast, whereas the other can be rescued without any interruption. Since the marginal harm to either engineer of having to wait to be rescued second would be the same, we decide to rescue first the one that requires no interruption. The harms to the TV viewers, though trumped, still tip the balance in this case of a tie. But now we learn that the other engineer is suffering very slightly more. He is trapped in a hotter area of the building, and he finds heat unpleasant. It may seem odd that this very slight change should cause us to reverse our decision. Although, strictly speaking, there is no tie between the engineers, it is 'close enough' to being a tie. If the harms to the TV viewers are significant enough to break a strict tie, it may seem, they must also be weighty enough to break what is very nearly a tie.

This may be put more formally in terms of the following condition.

Expanded Tie-breaking. For any x, y , with $0 < x < y$, if $(0, x, y) > (y, 0, 0)$ then for some $\varepsilon > 0$, $(0, x, y - \varepsilon) \geq (y, 0, 0)$.

This requires that if a less significant harm can tip the balance in a conflict between two equal more significant harms, then it can also do so in an arbitrarily close conflict. Lexical Generalized Collectivism and Extended Tie-Breaking together entail Extreme Collectivism.²³ Insofar as one accepts Transposition Pareto, it

²²The binary division of harms into 'relevant' and 'irrelevant' may seem too coarse-grained in some cases. It could be the case that, among the subset of irrelevant harms, some of these trump others. In this case, if there is a tie between relevant harms, it may seem inappropriate simply to aggregate all of the irrelevant harms, rather than giving priority to preventing the trumping irrelevant harm. This problem may be solved by introducing a hierarchy of relevant harms: 'Level 1 harms', 'Level 2 harms', and so on. For a formal definition see the [Appendix](#).

²³Lexical Generalized Collectivism implies $(0, \rho, 2) > (2, 0, 0)$. If $\rho > 0$, then $\rho/2 < \rho$, and thus Lexical Generalized Collectivism implies $(2, 0, 0) > (0, \rho, 2 - \varepsilon)$, for any $\varepsilon > 0$, thereby violating Extended Tie-Breaking.

seems very hard to reject Extended Tie-Breaking. If one allows irrelevant harms to break ties, then there seems little justification for not allowing them also to break almost-ties, especially given that an almost-tie can be arbitrarily close to a tie. Therefore, in my view, Lexical Generalized Collectivism also fails as an interpretation of the Close Enough View.

3.3. Numbers

I now consider a different approach to generalizing Extreme Collectivism. I shall argue that, though this approach has some advantages over the previous, it ultimately also fails to provide a satisfactory interpretation of the Close Enough View.

As noted above, the issue of aggregation is often put in terms of whether ‘numbers should count’. On an aggregationist view, we do allow numbers to count, in the sense that we allow the fact that a less significant harm will be suffered by a larger number of individuals to count in favour of preventing this harm. On a non-aggregationist view, we do not allow numbers to count in this way. On the approach I will consider here, we achieve a moderate position by saying that the importance of numbers – the extent to which they count – decreases as they grow larger. Numbers are important when they are small, but as they grow larger, their importance steadily diminishes.

Formally, this idea may be implemented by the following condition.

Weighted Numbers. For some strictly increasing, concave, positive-valued function g ,

$a \geq b$ if and only if

$$g(|H_{ab}|)M_{i \in H_{ab}} f(b_i) - f(a_i) \leq g(|H_{ba}|)M_{i \in H_{ba}} f(a_i) - f(b_i)$$

Here M denotes the arithmetic mean.²⁴ So Weighted Numbers tells us to minimize the weighted average significance of unprevented harms, where the weighting is given by the function g . This determines the extent to which numbers count as they increase. If g is linear, then numbers always count, and Weighted Numbers collapses into Extreme Collectivism. However, if g is strictly concave, then numbers count less as they increase. Moreover, if g is bounded above then Weighted Numbers becomes consistent with Moderate Non-Aggregation and Moderate Aggregation.

Suppose $x < x^+$ and $y < y^+$. Then Weighted Numbers implies that $(x^+, n \times y) > (x, n \times y^+)$ if and only if

$$g(1)[f(x^+) - f(x)] > g(n)[f(y^+) - f(y)]$$

This holds for all values of n if and only if

$$g(1)[f(x^+) - f(x)] > \lim_{n \rightarrow \infty} g(n)[f(y^+) - f(y)]$$

Now consider, for example, the function $g(x) = x/(x + 1)$. Notice, this function is bounded above. In particular, we have $\lim_{n \rightarrow \infty} g(n) = 1$. Thus, given this function, there exists an n such that $(x^+, n \times y) \not> (x, n \times y^+)$ if and only if

$$\frac{1}{2}[f(x^+) - f(x)] \leq [f(y^+) - f(y)]$$

²⁴That is $M_{i \in X} \phi(i) = \frac{1}{|X|} \sum_{i \in X} \phi(i)$.

So $\rho = 1/2$, and therefore this satisfies both Moderate Non-Aggregation and Moderate Aggregation.²⁵

Weighted Numbers is compatible with the conjunction of Moderate Aggregation, Moderate Non-Aggregation, Transposition Pareto and Extended Tie-Breaking. So it may appear to be an improvement over Generalized Collectivism and Lexical Generalized Collectivism. However, Weighted Numbers also has the following very problematic feature. Say that the harm to i in a ‘weakly trumps’ the harm to j in b , according to γ , if and only if there exists some n such that, for any m ,

$$(n \times b_i, m \times b_j) \succ (n \times a_i, m \times a_j).$$

Suppose the harm to i in a is only slightly greater than the harm to j in b . These harms might be, for example, suffering a mild headache for 10 minutes and 9 minutes, respectively. Then to say that the former harm weakly trumps the latter is to say that, for some number n , preventing n 10-minute headaches is better than preventing *any number* of 9-minute headaches. This is, of course, an absurd thing to say. However, this is exactly what is implied by a moderate form of Weighted Numbers. This condition combined with Moderate Non-Aggregation implies that a more significant harm always weakly trumps a less significant harm, regardless of how small the difference in significance may be.²⁶ On Weighted Numbers, we cannot have any trumping at all unless we have ubiquitous weak trumping. This is not quite extreme non-aggregationsim, but it is close enough to disqualify Weighted Numbers as an interpretation of the Close Enough View.

3.4. Individualism

I turn now to the individualist approach. Again, I begin by defining an extreme form, before considering how this may be generalized to allow moderation.

The extreme form of individualism tells us to minimize the significance of the most significant unprevented harm. This gives us the following condition.²⁷

Extreme Individualism. $a \geq b$ if and only if

$$\max_{i \in H_{ab}} f(b_i) - f(a_i) \leq \max_{i \in H_{ba}} f(a_i) - f(b_i)$$

Extreme Individualism is similar to Generalized Collectivism, except it counts as ‘irrelevant’ every harm other than the most significant.²⁸ Thus Extreme Individualism

²⁵This sort of moderate view was suggested to me by John Broome. It is similar to a view discussed, in a different context, by Binmore and Voorhoeve (2003).

²⁶Suppose g is bounded above (as required for consistency with Moderate Non-Aggregation). Then if $f(x^+) - f(x) > f(y^+) - f(y)$, there must be some n such that

$$g(n)[f(x^+) - f(x)] > \lim_{m \rightarrow \infty} g(m)[f(y^+) - f(y)]$$

²⁷Extreme Individualism is isomorphic to the so-called ‘minimax regret rule’ for individual choice. See e.g. Binmore (2009: 156–157).

²⁸One might suspect that Generalized Collectivism with $\rho = 1$ would be equivalent to Extreme Individualism. But this is not quite so. Suppose the most significant harms in a and b are equally

is also inconsistent with Transposition Pareto. But here the problem is more severe. Consider the following outcomes.

	1	2	3
<i>a</i>	1	1	0
<i>b</i>	0	0	1

In this case, Transposition Pareto implies $a \succ b$, but Extreme Individualism implies $a \sim b$, because the most significant harm is the same in both outcomes. It seems wrong, however, to allow harm to befall two people in order to prevent the same harm befalling only one person. Even some extreme non-aggregationists may find this too extreme.²⁹

This problem may again be solved by adopting a lexical variant. Here I define this condition only informally (see [Appendix](#) for a formal definition).

Lexical Extreme Individualism.

1. Choose so that either
 - (a) there are no unprevented harms, or
 - (b) the most significant unprevented harm is no more significant than the most significant prevented harm.
2. In the event of a tie at step 1, choose so that either
 - (a) there is only one unprevented harm and at least two prevented harms, or
 - (b) the second-most significant unprevented harm is no more significant than the second-most significant prevented harm.
3. And so on.

Note that in this definition, for example, the ‘second-most significant’ harm may be equally significant as the most significant.³⁰

Clearly Extreme Individualism and Lexical Extreme Individualism are both incompatible with Moderate Aggregation. If the most significant harm in one outcome is more significant than that in the other, then no number of less significant harms, no matter how many, can outweigh this. How might these conditions be generalized so as to allow moderation? Here is one idea. Extreme Individualism requires that the most significant prevented harm must be at least as significant as every unprevented harm. We may weaken this by requiring only that the most significant prevented harm is not *trumped* by any unprevented harm. This gives us the following condition.

significant, but two individuals suffer this harm in *a*, whereas only one suffers it in *b*. Then Generalized Collectivism with $\rho = 1$ prefers *b*, but Extreme Individualism is indifferent.

²⁹One person who famously endorses such extremism is Taurek (1977).

³⁰More precisely, let *h* be a bijection from $\{1, 2, \dots, |H_{ab}|\}$ into H_{ab} such that $i < j$ implies $f(b_{h(j)}) - f(a_{h(j)}) \leq f(b_{h(i)}) - f(a_{h(i)})$. Then the *i*-th most significant harm in *a* is the harm to *h(i)*.

Generalized Individualism. $a \geq b$ if and only if

$$\rho \max_{i \in H_{ab}} f(b_i) - f(a_i) \leq \max_{i \in H_{ba}} f(a_i) - f(b_i)$$

This definition is identical to that of Extreme Individualism except for the insertion of the ratio ρ . Thus when $\rho = 1$, these two conditions become equivalent. The idea is that moderation is achieved by setting $\rho < 1$. Then this condition will require only that we come ‘close enough’ to minimizing the maximum significance of unprevented harms.

The same idea may be applied to Lexical Extreme Individualism.

Lexical Generalized Individualism. The same as Lexical Extreme Individualism, except that ‘is no more significant than’ is replaced by ‘does not trump’.

The problem with this, however, is that, given Regularity, it is in fact impossible that $\rho < 1$. So these generalized conditions simply collapse into their extreme forms. For any $x > y > 0$, Regularity implies $(x, 0) > (0, y)$. It follows, given either Generalized Individualism or Lexical Generalized Individualism, that for any $x > y > 0$,

$$f(y) - f(0) < \rho(f(x) - f(0))$$

But this implies $\rho = 1$. So these generalized conditions do not allow moderation after all, at least not given Regularity.

To give a concrete example, suppose we can either prevent a headache to one person, or prevent a slightly less intense headache to another person. Assume also that the wellbeing level either individual will have if she does not suffer the headache is the same. Since the difference in harm is only slight, the greater harm does not trump the smaller harm, and therefore both Lexical Extreme Individualism and Lexical Generalized Individualism are indifferent between these options. Regularity, on the other hand, requires that we prevent the worse headache.³¹

This problem is solved by the following condition.³²

Individualist-Collectivist Hybrid.

1. If Lexical Generalized Individualism selects a uniquely permissible option, then choose this option.
2. Otherwise choose according to Extreme Collectivism.

This is a generalization of both Extreme Collectivism and Lexical Extreme Individualism. If $\rho = 0$, then no harm trumps any other harm. In this case, Lexical Generalized Individualism will always imply that both options are permissible, so the decision will be made by Extreme Collectivism. On the other hand, if $\rho = 1$,

³¹Doubts about this implication of Regularity may be met, I believe, by the points I made earlier regarding vagueness and degrees of wrongness. We may feel that, in practice, it is not possible to know that one headache is only slightly less intense. Also, Regularity does not imply that failing to prevent the worse headache would be *seriously* wrong; it may be only slightly wrong.

³²A view like this was suggested to me by Christopher Jay.

then Lexical Generalized Individualism becomes equivalent to Lexical Extreme Individualism, which judges both options to be permissible only if the harms in each are exactly the same. In this case, Extreme Collectivism agrees with Lexical Extreme Individualism. Therefore, Individualist-Collectivist Hybrid is equivalent to Extreme Collectivism when $\rho = 0$, and equivalent to Lexical Extreme Individualism when $\rho = 1$.

Individualist-Collectivist Hybrid is compatible with Regularity even with $\rho < 1$. In the example of the two headaches above, Lexical Extreme Individualism fails to select a uniquely permissible option, and so Extreme Collectivism decides this case in favour of preventing the worse headache. Thus Individualist-Collectivist Hybrid does not collapse into extremism. Moreover, it is consistent with both Transposition Pareto and Extended Tie-Breaking. So in these respects, it is preferable to the (purely) collectivist conditions

I suspect the main criticism of Individualist-Collectivist Hybrid will be that it seems rather *ad hoc*. It lacks a unified rationale for the verdicts it gives, sometimes appealing to individualist considerations, while at other times appealing to quite different collectivist considerations. However, some degree of disunification seems inevitable on a moderate view. Indeed, as the earlier quotation from Scanlon makes clear, it is the whole point of the Close Enough View to distinguish two classes of cases, and to treat these in different ways. I conclude, therefore, that among the candidates surveyed above, the most promising interpretation of the Close Enough View, in the context of binary choices only, is Individualist-Collectivist Hybrid.

4. Non-binary choices

I turn now to non-binary choices. As I shall show, problems emerge here for *all* (minimally plausible) moderate views, because these views entail the existence of ‘obligation cycles’. As a result, these views must violate certain consistency conditions in the context of non-binary choice. I first explain how moderation leads to obligation cycles, and why this might be problematic, before considering some possible solutions.

4.1. Obligation cycles

An important consequence of moderation is that the ‘close enough’ relation must be intransitive. To illustrate, suppose f is linear and $\rho = 1/2$. Then, for example, a harm of size 1 is close enough to a harm of size 2, which is in turn close enough to a harm of size 3. But the first of these harms is not close enough to the third.³³ This of course

³³That this holds in general, for any moderate choice function, can be proven as follows. Since f is continuous (because concave) and strictly increasing, there exist $x, y \in \mathbb{R}_+$ such that

$$\begin{aligned} f(x) &= \rho f(1) + (1 - \rho)f(0) \\ f(y) &= \rho^2 f(1) + (1 - \rho^2)f(0) \end{aligned}$$

Then we have

$$\frac{f(x) - f(0)}{f(1) - f(0)} = \rho \quad \frac{f(y) - f(0)}{f(x) - f(0)} = \rho \quad \frac{f(y) - f(0)}{f(1) - f(0)} = \rho^2$$

fits with the ordinary meaning of ‘close enough’. Suppose, for example, a married couple is contemplating taking jobs in different cities, but want to ensure that they will still be ‘close enough’ to each other. They may feel that, say, London is close enough to Paris, and that Paris is close enough to Geneva, but that London is not close enough to Geneva.

As others have observed, this intransitivity can have troubling consequences.³⁴ Consider, for example, applying Generalized Collectivism to the following outcomes.

	1	2	3	4	5	6
<i>a</i>	0	4	4	3	3	3
<i>b</i>	7	0	0	3	3	3
<i>c</i>	7	4	4	0	0	0

Assume f is linear and $\rho = 1/2$. Compare first a and b . Here the two harms of size 4 in b (to individuals 2 and 3) are close enough to the one harm of size 7 in a (to individual 1). So, because all harms are relevant, and the total harm is less in a ($7 < 4+4$), we have $a > b$. Likewise, comparing b and c , because all harms are relevant, and the total harm is less in b , we have $b > c$. But the situation is different when comparing a and c . In this case, the harms in c are trumped by the harm in a , and therefore are declared irrelevant. So we have $c > a$. In summary, we have a cycle: $a > b$, $b > c$, $c > a$.³⁵ I call this an ‘obligation cycle’ because, recall, $a > b$ means that it is only permissible to choose a (when a and b are the only options), or, in other words, that it is obligatory to choose a .

A similar set of outcomes can be constructed for any f and ρ , provided only that $0 < \rho < 1$. Thus the combination of Generalized Collectivism, Moderate Aggregation and Moderate Non-Aggregation is inconsistent with the following condition.

No Obligation Cycles. For any sequence of outcomes a^1, a^2, \dots, a^n , if $a^i > a^{i+1}$ for all i , then $a^n \not> a^1$.

In this above example, the cycle is a consequence of moderation, resulting from the combination of both Moderate Non-Aggregation and Moderate Aggregation. In the first two comparisons (a vs b and b vs c), the aggregationist part of the view is employed, whereas in the third comparison (a vs c), the non-aggregationist part is employed. It is this combination of aggregationism and non-aggregationism which, predictably, gives rise to a cycle. The cycle may be avoided, by ‘going extreme’ in either direction. In the above example, Extreme Collectivism implies $a > b > c$, whereas Extreme Individualism implies $c > b > a$.

In other cases, however, obligation cycles may arise without *any* aggregation. Moderate Non-Aggregation, Non-Arbitrariness and No Obligation Cycles are

Now if $0 < \rho < 1$, then $\rho^2 < \rho$. So in this case, the harm of getting 0 instead of y is close enough to the harm of getting 0 instead of x , which is in turn close enough to the harm of getting 0 instead of 1. But the first of these harms is not close enough to the third.

³⁴See e.g. Norcross (2002), Parfit (2003), Fleurbaey *et al.* (2009) and Voorhoeve (2014).

³⁵This is also implied by Lexical Generalized Collectivism and Individualist-Collectivist Hybrid.

jointly inconsistent. (Notice, Moderate Aggregation is not included here.) The proof is similar to that given above in the context of Permutation Invariance. Without loss of generality, assume as before that

$$f(1/2) - f(0) < \rho(f(1) - f(0))$$

Now consider the following outcomes.

	1	2	3
<i>a</i>	0	1/2	1
<i>b</i>	1/2	1	0
<i>c</i>	1	0	1/2

Compare first *a* and *b*. As shown above, it follows that the one harm in *b* trumps both harms in *a*, and so, by Non-Arbitrariness, we have $a > b$. But by the same argument we also have $b > c$ and $c > a$. So we have an obligation cycle.

Obligation cycles raise special problems in the context of non-binary choice. Suppose we have an obligation cycle involving a number of options. What then should we say about a choice in which *all* of these options are available? According to Parfit, we must then say that *none* of the options is permissible (Parfit 2003: 384). That is, we must reject the following condition.

No Prohibition Dilemmas. $\gamma(A) \neq \emptyset$.

But this is not quite correct. Rather, we must reject *either* No Prohibition Dilemmas or the following condition.

Contraction Consistency. Let *A* be a subset of *B*. Then, for any $a \in A$, $a \in \gamma(B)$ only if $a \in \gamma(A)$.

This requires that a permissible option cannot become impermissible merely as a result of removing some other options. Contraction may, of course, have the opposite effect: transforming an option from impermissible to permissible. This might happen if some superior option is removed. But surely, one might think, if an option is good enough to choose from the larger set of options, then it must also be good enough to choose from the smaller set. Thinning the competition must, if anything, make it easier to win, not harder.

No Prohibition Dilemmas and Contraction Consistency together entail that $>$ is acyclical. Thus (assuming Non-Arbitrariness) Moderate Non-Aggregation forces us to reject one of these two conditions. One simple way to extend a choice function from binary to non-binary choices is by the following condition.

Maximization. For any $a \in A$, $a \in \gamma(A)$ if and only if, for any $b \in A$, $a \geq b$.

This says that an option is permissible just in case it is maximal relative to the relation \geq . If this relation is interpreted as saying that one option is 'at least as good'

as another, then Maximization tells us to choose one of the ‘best’ options. Clearly Maximization implies Contraction Consistency. So if $>$ is cyclical, then Maximization is incompatible with No Prohibition Dilemmas. (If every option is worse than some other option, then no option is best.)

However, there are alternatives to Maximization which are compatible with No Prohibition Dilemmas, even when $>$ is cyclical. One such alternative substitutes \geq with its transitive closure \geq^+ . This gives us:

Transitive Closure Maximization. Let \geq^+ be the transitive closure of \geq restricted to A .³⁶ Then, for any $a \in A$, $a \in \gamma(A)$ if and only if, for any $b \in A$, $a \geq^+ b$.

If we have, for example, $a > b$, $b > c$, $c > a$ (an obligation cycle), then Transitive Closure Maximization implies $\gamma(\{a, b, c\}) = \{a, b, c\}$. An intuitive rationale for this view is that the three options are morally ‘on a par’: each is worse than some alternative. As there is nothing to separate the options, we may conclude that all three are permissible. Provided \geq is complete (which amounts to saying there are no prohibition dilemmas in binary choices), Transitive Closure Maximization implies No Prohibition Dilemmas. So if $>$ is cyclical, then Transitive Closure Maximization is incompatible with Contraction Consistency.

Voorhoeve suggests another way of extending Generalized Collectivism. In a non-binary choice, determining the total relevant harm in a given option is more complicated. This is because, as noted above, determining the counterfactual harm to an individual is more complicated. Which of the alternatives to the chosen option should be regarded as what would have happened otherwise? One solution is to identify this with the *best* that could have happened. That is, the harm to an individual in an outcome is determined by comparing her wellbeing in this outcome with her *maximum* wellbeing in any outcome. Formally, for an outcome $a \in A$, individual i is harmed in a , relative to A , if and only if $a_i < \max_{b \in A} b_i$.

The set of individuals harmed in a , relative to A , is $H_{aA} = \bigcup_{b \in A} H_{ab}$. And the set of individuals *relevantly* harmed in a , relative to A , may then be defined as

$$R_{aA} = \{i \in H_{aA} : \forall b \in A \setminus \{a\} \forall j \in I_a \setminus \{i\} f(a_i) - f(\max_{c \in A} c_i) \geq \rho(f(b_j) - f(\max_{c \in A} c_j))\}.$$

This gives us the following condition.

Global Generalized Collectivism. For any $a \in A$, $a \in \gamma(A)$ if and only if, for any $b \in A$,

$$\sum_{i \in R_{aA}} \max_{c \in A} f(c_i) - f(a_i) \leq \sum_{i \in R_{bA}} \max_{c \in A} f(c_i) - f(b_i)$$

³⁶More precisely, \geq^+ is the smallest transitive relation on A such that, for all $a, b \in A$, if $a \geq b$ then $a \geq^+ b$. Conditions like this are discussed in Sen (1977), for example.

This clearly implies Non-Arbitrariness and No Prohibition Dilemmas. Therefore, in combination with Moderate Non-Aggregation it violates Contraction Consistency.

To summarize, moderation entails the existence of obligation cycles (assuming the very plausible Non-Arbitrariness). Therefore, moderates must reject either No Prohibition Dilemmas or Contraction Consistency, but they have a choice as to which they reject. If they prefer to reject No Prohibition Dilemmas, they may adopt Maximization. On the other hand, if they prefer to reject Contraction Consistency, then they have at least two options. One is to replace Maximization with Transitive Closure Maximization. Another is to adopt a view like Global Generalized Collectivism.

In my view, Non-Arbitrariness and No Prohibition Dilemmas are non-negotiable. Thus, in the remainder of the paper, I consider the prospects of rejecting Contraction Consistency.

4.2. *Companions in guilt*

Are violations of Contraction Consistency a cost worth paying for moderation? Contraction Consistency is not uncontroversial. As others have shown, this condition is violated by various well-established and seemingly consistent views.³⁷ Moderate aggregationists may therefore mount a ‘companions in guilt’ defence. If their view is convicted of inconsistency for violating Contraction Consistency, then many others must be convicted too. However, this defence, I argue, is not entirely convincing. Moderate views differ from these other views in an important way, so they should not be prosecuted as co-defendants.

One class of views which violate Contraction Consistency involve ‘satisficing’. Unlike maximizing views, which insist on choosing the best option, satisficing views permit choosing a non-best option, provided it is at least ‘good enough’. This may seem to cohere more with so-called ‘commonsense morality’. For example, while many accept that they ought to give to charity, few feel obligated to give as much as maximizing seems to demand. From a maximizing perspective, Contraction Consistency makes perfect sense; from a satisficing perspective, less so. Consider, for example, a simple satisficing view which merely forbids choosing the worst option, permitting all else. Given a choice between three options, Good, Mediocre and Bad, listed from best to worst, this view permits choosing either Good or Mediocre. But when Bad is removed, Mediocre becomes the worst and therefore no longer permissible, thereby violating Contraction Consistency.

Another well-known example, suggested by Sen (1993: 501), has a similar structure. Imagine three slices of cake of varying size: Small, Medium and Large. Etiquette dictates that when offered a selection of slices of cake, one ought to forgo the largest slice, because to take it would be greedy. But this implies that choosing Medium is permissible when all three slices are offered, but not when Large is withdrawn.

These views seem at least *consistent*. As Sen points out, the label ‘Contraction Consistency’ is really a misnomer (Sen 1993: 499). This condition is not merely

³⁷See e.g. Sugden (1985), Sen (1993) and Tungodden and Vallentyne (2005).

a requirement of consistency, at least not in the strict logical sense. The following statements are not logically inconsistent:

1. Choosing a from $\{a, b, c\}$ is permissible.
2. Choosing a from $\{a, b\}$ is not impermissible.

Inconsistency arises only with the addition of certain background assumptions. For example, these statements are jointly inconsistent with the following:

3. Choosing an option from a set of options is permissible if and only if this option is not worse than any option in this set.

Together (1) and (3) entail that a is not worse than b . But (2) and (3) entail that a is worse than b .³⁸ But this is unlikely to trouble proponents of satisficing, because they reject the background assumption. On their view, an option can be permissible even when it is worse than some other option.

Notice, however, that satisficers need not reject (3) in its entirety. They need only reject the left-to-right part, not the right-to-left. That is, while satisficers deny that being the best option (or one of the best) is *necessary* for being permissible, they accept that this is *sufficient*. However, the conjunction of Moderate Non-Aggregation and Non-Arbitrariness is incompatible with both necessity and sufficiency, unless the 'better than' relation is cyclical. As shown above, these conditions together entail obligation cycles. Given acyclicity of 'better than', this implies that some option is impermissible even though it is not worse than any alternative. Notice, neither the satisficing nor etiquette view entails obligation cycles. Thus a companions in guilt defence seems unconvincing. In the case of these other views, there is a justification for violating Contraction Consistency which is not also available to moderate views.

4.3. Intrinsicism

Voorhoeve pursues a different response. His argument may be presented as follows.³⁹ Violations of Contraction Consistency reveal genuine inconsistencies only in cases where contracting the choice situation does not have the effect of altering any relevant properties of the remaining options. If only the *intrinsic* properties of an option are relevant, where this excludes, in particular, the option's relations to other options, then *all* violations reveal genuine inconsistencies, because intrinsic properties are preserved by contraction. According to some views, however, an option's *extrinsic* properties may also be relevant. Therefore, such a view may violate Contraction Consistency without being inconsistent at all. Now, one might object that such views mistakenly count as relevant properties that really are irrelevant. But the objection, in this case, would not be that the view is *inconsistent*. Given this view's account of

³⁸This assumes that 'worse than' is irreflexive: an option cannot be worse than itself.

³⁹This is not exactly how Voorhoeve presents the argument. But I believe it amounts to the same thing.

relevant properties, its prescriptions in contraction cases may be perfectly consistent.

For example, it may be argued as follows that the etiquette view is not really inconsistent. Taking the medium slice when both the large and the small slices are also on offer differs in a relevant way from doing this when the only alternative is the small slice: in the latter case, but not the former, taking the medium slice displays greed or a lack of manners. Therefore, taking the medium slice is not *relevantly* the same in both cases.⁴⁰ Of course, the medium slice of cake itself does not change in its intrinsic properties. Rather, what changes is how it compares to the other slices on offer. In the second situation, but not in the first, it is the largest slice. Therefore, since the etiquette view counts this extrinsic property as relevant, there is nothing inconsistent in its judgement that choosing the medium slice is permissible in the first situation, but not in the second.

My framework may be expanded to reflect these ideas as follows. Let an 'extended option' be a pair (a, A) where $a \in A$. For neatness, I shall write a_A instead of (a, A) . We may think of a_A as 'choosing a from A '. Now, let a distributive view be represented by a choice function γ combined with an equivalence relation \equiv on the set of extended options. Interpret $a_A \equiv b_B$ as meaning that all the relevant properties of a in A are the same as those of b in B .⁴¹

We may then define the following condition.

General Contraction Consistency. Let A be a subset of B such that, for any $a \in A$, $a_A \equiv b_B$. Then, for any $a \in A$, $a \in \gamma(B)$ only if $a \in \gamma(A)$.

This may be regarded as the result of weakening Contraction Consistency so that violations occur only in cases where all relevant properties are preserved by contraction. A view that violates this condition seems genuinely inconsistent.

Now, General Contraction Consistency collapses into Contraction Consistency given the following condition.

Intrinsicalism. If $a = b$ then $a_A \equiv b_B$.

This says in effect that the relevant properties of an option do not include its relations to other options, so these remain the same in every choice situation.

However, moderates may reject Intrinsicalism. Consider, for example, Global Generalized Collectivism. This may be made compatible with General Contraction Consistency by adopting an equivalence relation such that $a_A \equiv b_B$ if and only if $a = b$ and, for all $i \in I_a$, $\max_{c \in A} c_i - a_i = \max_{c \in B} c_i - b_i$. Obviously this violates Intrinsicalism. But I can think of no weighty independent reason to accept this condition. If we find a moderate view otherwise attractive, then the fact that it must reject Intrinsicalism in order to be consistent does not seem a strong objection. Perhaps one consideration in favour of Intrinsicalism is

⁴⁰This general strategy of refining the description of options is explored in detail by Broome (1991: Ch. 5).

⁴¹We could instead represent a distributive view as a choice function plus a set of relevant properties, and then derive the equivalence relation from the latter. For simplicity, I skip this intermediate step.

practical convenience. Decision making is easier if we can treat options as relevantly the same in all choice situations. This consideration, however, seems fairly easily outweighed.

One might worry that General Contraction Consistency is almost entirely vacuous. Any choice function, one might suspect, can be combined with an equivalence relation that satisfies General Contraction Consistency. We might add that it is not enough merely to select any old equivalence relation; there must be some story to tell justifying this selection. But unless a view is entirely arbitrary or random, there will surely be some such story to tell. What this shows, I believe, is that consistency is a very weak requirement. We will not get very far in evaluating competing views merely by considerations of consistency. We need to engage with the *substance* of these views.

4.4. Local vs global

I've argued that we need not be troubled by violations of Contraction Consistency. But this still leaves two different approaches to extending a moderate choice function from binary to non-binary choices. Focus now on Generalized Collectivism. One way to extend this to non-binary choices is by combining it with Transitive Closure Maximization. I call this the 'local approach'. The other way is to adopt Global Generalized Collectivism. I call this the 'global approach'. Both approaches satisfy No Prohibition Dilemmas but violate Contraction Consistency (assuming Moderate Non-Aggregation). Which is preferable? I argue for the local approach.

Notice first that these approaches diverge only in 'non-symmetrical' cases. By a 'symmetrical' case I mean one in which the options are generated by successive applications of the same permutation. An example like this is used above to show that Moderate Non-Aggregation and Non-Arbitrariness alone are sufficient to create an obligation cycle, without the aid of Moderate Aggregation. In symmetrical cases, both approaches imply that *all* options are permissible. This is surely the right conclusion. Due to the symmetry, there really is nothing to decide between the options, and so it would be entirely arbitrary to say one option is permissible but another is not. The only other non-arbitrary conclusion is that no option is permissible (a prohibition dilemma), but this too would be absurd. In non-symmetrical cases, the local approach still implies that all options are permissible, but not the global approach.

Consider again the non-symmetrical example presented above.

	1	2	3	4	5	6
<i>a</i>	0	4	4	3	3	3
<i>b</i>	7	0	0	3	3	3
<i>c</i>	7	4	4	0	0	0

Assume as before that f is linear and $\rho = 1/2$. Then, as we've seen, Generalized Collectivism implies an obligation cycle: $a > b$, $b > c$, $c > a$. Combining this with

Transitive Closure Maximization, we get $\gamma(\{a, b, c\}) = \{a, b, c\}$. On the other hand, Global Generalized Collectivism implies $\gamma(\{a, b, c\}) = \{c\}$.

But the latter implication seems dubious. In our deliberation between these options, we may entertain various hypothetical questions. For example, we may ask: 'Assuming for now that it is permissible to choose c , is it also permissible to choose b ?' Now, the harms in c are, we have assumed, close enough to those in b , and moreover the total harm in c is greater than that in b . So we should conclude that, if it is permissible to choose c , then it is also permissible to choose b . It would be illegitimate to disregard the harms in c on the grounds that these are trumped by the harms in a . The question we are answering is about the permissibility of b and c , not a . When answering this hypothetical question, the only relevant considerations concern how b compares to c ; it is irrelevant how either of these options compares to a . Therefore, the question reduces to the binary choice $\{b, c\}$. If we believe that it is permissible to choose b from $\{b, c\}$, then we should conclude that, on the assumption that it is permissible to choose c from $\{a, b, c\}$, it is also permissible to choose b from $\{a, b, c\}$. However, Global Generalized Collectivism contradicts this conclusion, because it says only c is permissible.

To put this more formally, my objection to Global Generalized Collectivism is that (combined with Moderate Non-Aggregation) it violates the following condition.⁴²

Expansion Consistency. For any $a, b \in A$, If $a \geq b$ and $b \in \gamma(A)$, then $a \in \gamma(A)$.

This condition seems to me very plausible. As with Contraction Consistency, I do not consider Expansion Consistency to be a requirement of consistency, despite its name. So my objection is not that Global Generalized Collectivism is inconsistent. Moreover, I am not denying that some extrinsic properties of options may be relevant. It is true, for example, that the extrinsic properties of a change when the choice situation expands from $\{a, b\}$ to $\{a, b, c\}$. But these changes should not be relevant to the hypothetical question whether a is permissible on the assumption that b is permissible. With respect to *this* question, the relevant extrinsic properties of a are only its relations to b , and these do not change with the addition of c .

To further emphasize this point, consider the following example.

	1	2	3	4	5	6
a	0	0	0	0	0	0
b	7	0	0	3	3	3
c	7	4	4	0	0	0

⁴²This is weaker than the condition often called 'Expansion Consistency'. But it is strong enough for our purposes.

In this case, there is no obligation cycle. Generalized Collectivism implies $b > c > a$. Combining this with Transitive Closure Maximization, gives the intuitively correct result, $\gamma(\{a, b, c\}) = \{b\}$.⁴³ But Global Generalized Collectivism implies $\gamma(\{a, b, c\}) = \{c\}$, again violating Expansion Consistency. The harms in c are counted as irrelevant, because they are trumped by the maximum harm in a . But again this does not seem a legitimate reason to choose c ahead of b , since the harms in c are not trumped by any harm in b . One might say, a is so clearly unacceptable that it should not really be regarded as an option at all, making this in effect a choice between b and c only. And Global Generalized Collectivism could be modified in this way to avoid the problem. We could say, for example: first eliminate any Pareto-inferior options (in this case a), and then apply Global Generalized Collectivism to the remaining options. But this seems a little *ad hoc*, and, in any case, the problem with the previous example would still remain. A better solution is to adopt Transitive Closure Maximization.

5. Conclusion

This completes my survey of candidate interpretations of the Close Enough View. I conclude that the most promising of the candidates considered is the combination of Individualist-Collectivist Hybrid and Transitive Closure Maximization. I cannot claim that this is the best possible interpretation. My search was restricted to regular choice functions, and even within this subclass, there may be other options that I have not considered. Still, I hope to have shown at least that this interpretation is worthy of further investigation.

What does this tell us about the Close Enough View? Suppose my preferred candidate is in fact the best we can do here. Is this good enough? Should we then accept the Close Enough View, on this interpretation, or instead abandon moderation? Individualist-Collectivist Hybrid may have some unappealing features. It is not an elegant condition, and, as noted above, it may be considered *ad hoc*. I should also note that extending this condition to cases of uncertainty may be far from straightforward. It definitely cannot be anything so simple as maximizing the expected value of a social utility function. Moreover, as noted above, the case for moderation rests largely on intuitions about cases, and we have reason to think these intuitions may be unreliable. So those with a taste for simplicity, or a less firm commitment to moderation, might see the investigation above as an argument against the Close Enough View. In any case, I hope to have clarified the issue of aggregation, by displaying some of the options available to moderates, and some of the costs involved in accepting these.

Acknowledgements. Earlier versions of this paper were presented at the University of York, the London School of Economics Choice Group and the Oxford Moral Philosophy Seminar. I am thankful to participants at these events for their helpful comments. I am especially thankful to Alex Voorhoeve and two anonymous referees for their comments on written versions of the paper.

⁴³When \succeq is transitive, Transitive Closure Maximization agrees with Maximization.

References

- Binmore K.** 2009. *Rational Decisions*. Princeton, NJ: Princeton University Press.
- Binmore K. and A. Voorhoeve** 2003. Defending transitivity against Zeno’s paradox. *Philosophy and Public Affairs* **31**, 272–279.
- Broome J.** 1991. *Weighing Goods*. Oxford: Blackwell.
- Broome J.** 2004. *Weighing Lives*. Oxford: Oxford University Press.
- Brown C.** 2007. Prioritarianism for variable populations. *Philosophical Studies* **134**, 325–361.
- Fleurbaey M., B. Tungodden and P. Vallentyne** 2009. On the possibility of nonaggregative priority for the worst off. *Social Philosophy and Policy* **26**, 258–285.
- Kamm F.M.** 2005. Aggregation and two moral methods. *Utilitas* **17**, 1–23.
- Norcross A.** 2002. Contractualism and aggregation. *Social Theory Practice* **28**, 303–314.
- Parfit D.** 2003. Justifiability to each person. *Ratio* **16**, 368–390.
- Scanlon T.** 1998. *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Sen A.** 1977. Social choice theory: a re-examination. *Econometrica* **45**, 53–89.
- Sen A.** 1993. Internal consistency of choice. *Econometrica* **61**, 495–521.
- Sugden R.** 1985. Why be consistent? A critical analysis of consistency requirements in choice theory. *Economica* **52**, 167–183.
- Taurek J.M.** 1977. Should the numbers count? *Philosophy and Public Affairs* **6**, 293–316.
- Tungodden B. and P. Vallentyne** 2005. On the possibility of Paretian egalitarianism. *Journal of Philosophy* **102**, 126–154.
- Voorhoeve A.** 2014. How should we aggregate competing claims? *Ethics* **125**, 64–87.

Appendix

This appendix provides some formal definitions that are too complicated to be included in the main text. I first define an ‘extended’ version of Lexical Generalized Collectivism, which recognizes a more fine-grained division of options into levels of relevance. These levels may be recursively defined as follows.

1. $R_{ab}^1 = R_{ab}$
2. for $k > 1$,

$$R_{ab}^k = \{i \in H_{ab} : \forall j \in H_{ba} \setminus R_{ba}^{k-1} f(b_i) - f(a_i) \geq \rho(f(a_i) - f(b_j))\}$$

It may be helpful to think of these levels as being defined by the following procedure. Begin with H_{ab} and H_{ba} , the sets of individuals harmed in a and b respectively. Then define R_{ab}^1 and R_{ba}^1 exactly as R_{ab} and R_{ba} above. Next, set aside the individuals in R_{ab}^1 and R_{ba}^1 , and define R_{ab}^2 and R_{ba}^2 exactly as R_{ab}^1 and R_{ba}^1 , except this time consider only the remaining individuals. Then repeat this procedure to define R_{ab}^3 and R_{ba}^3 , and so on. (Notice, these levels are by definition ‘nested’: $R_{ab}^1 \subseteq R_{ab}^2 \subseteq R_{ab}^3 \dots$)

The idea is that higher levels (where R^1 is the highest) are given lexical priority over lower levels. First, we minimize the total significance of unprevented harms to individuals in R_{ab}^1 and R_{ba}^1 . In the case of a tie, we then minimize the same for individuals in R_{ab}^2 and R_{ba}^2 . And so on. This gives us the following condition.

Extended Lexical Generalized $a \geq b$ if and only if, for all k , if

$$\sum_{i \in R_{ab}^k} f(b_i) - f(a_i) > \sum_{i \in R_{ba}^k} f(a_i) - f(b_i)$$

then there exists some $j < k$ such that

$$\sum_{i \in R_{ab}^j} f(b_i) - f(a_i) < \sum_{i \in R_{ba}^j} f(a_i) - f(b_i).$$

Now I give a formal definition of Lexical Extreme Individualism. For outcomes a and b (with $\mathbb{N}_a = \mathbb{N}_b$), let h_{ab} be an injective function from $\{1, 2, \dots, |H_{ab}|\}$ into I_a such that, for any I, j , if $i < j$ then $f(b_{h_{ab}(i)}) - f(a_{h_{ab}(i)}) \geq f(b_{h_{ab}(j)}) - f(a_{h_{ab}(j)})$. So $h_{ab}(i)$ is the individual who suffers the i -th most

significant harm in a . Note that adjacent individuals in this sequence may suffer equally significant harms. Then we have the following definition.

Lexical Extreme Individualism.

1. If $|H_{ab}| \leq |H_{ba}|$, then $a \geq b$ if and only if, for all $1 \leq i \leq |H_{ba}|$, if

$$f(b_{h_{ab}(i)}) - f(a_{h_{ab}(i)}) > f(a_{h_{ba}(i)}) - f(b_{h_{ba}(i)})$$

then for some $1 \leq j < i$,

$$f(a_{h_{ba}(j)}) - f(b_{h_{ba}(j)}) > f(b_{h_{ab}(j)}) - f(a_{h_{ab}(j)})$$

2. If $|H_{ab}| < |H_{ba}|$, then $a > b$ if $a \geq b$, and $b > a$ if $a \not\geq b$.

A formal definition of Lexical Generalized Individualism is obtained by inserting ρ into the definition above in the appropriate places.

Campbell Brown is an Assistant Professor in the Department of Philosophy, Logic and Scientific Method at the London School of Economics. His research is primarily in moral and political philosophy, especially formal approaches to these.