

ARTICLES

# Lives, Limbs, and Liver Spots: The Threshold Approach to Limited Aggregation

S. Matthew Liao<sup>1</sup>  and James Edgar Lim<sup>2</sup>

<sup>1</sup>Center for Bioethics, New York University, New York, NY, USA and <sup>2</sup>School of Philosophy, Australian National University, Canberra, ACT, Australia

**Corresponding author:** S. Matthew Liao; Email: [matthew.liao@nyu.edu](mailto:matthew.liao@nyu.edu)

## Abstract

Limited Aggregation is the view that when there are competing moral claims that demand our attention, we should sometimes satisfy the largest aggregate of claims, depending on the strength of the claims in question. In recent years, philosophers such as Patrick Tomlin and Alastair Norcross have argued that Limited Aggregation violates a number of rational choice principles such as Transitivity, Separability, and Contraction Consistency. Current versions of Limited Aggregation are what may be called Comparative Approaches because they involve assessing the relative strengths of various claims. In this paper, we offer a non-comparative version of Limited Aggregation, what we call the Threshold Approach. It states that there is a non-relative threshold that separates various claims. We demonstrate that the Threshold Approach does not violate rational choice principles such as Transitivity, Separability, and Contraction Consistency, and we show that potential concerns regarding such a view are surmountable.

**Keywords:** limited aggregation; threshold; transitivity; nonconsequentialism

## 1. Introduction

Consider the following cases:

*Death vs. Liver Spots:* We can either save Ann, a young woman, from death, allowing her to live a normal healthy life for many years, or prevent some number of different people from having liver spots for the rest of their lives.

*Death vs. Loss of Limbs:* We can either save Bob, a young man, from death or prevent some number of different people from losing both their legs.<sup>1</sup>

<sup>1</sup>Similar pairs of cases have been mentioned by Alastair Norcross, “Comparing Harms: Headaches and Human Lives,” *Philosophy & Public Affairs* 26, no. 2 (1997): pp. 135–67; F. M. Kamm, “Aggregation and Two Moral Methods,” *Utilitas* 17, no. 1 (2005): pp. 1–23; Alex Voorhoeve, “How Should We Aggregate Competing Claims?,” *Ethics* 125, no. 1 (2014): pp. 64–87; Patrick Tomlin, “On Limited Aggregation,” *Philosophy & Public Affairs* 45, no. 3 (2017): pp. 232–60; Joe Horton, “Always Aggregate,” *Philosophy & Public Affairs* 46, no. 2 (2018): pp. 160–74; and Korbinian R uger, “Aggregation with Constraints,” *Utilitas* 32, no. 4 (2020): pp. 454–71.

  The Author(s), 2024. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

It seems that we should save Ann in *Death vs. Liver Spots*, regardless of how many people we can potentially save from liver spots. A similar policy applies for many other minor afflictions; we should save Ann instead of saving any number of people from broken fingernails, itchy noses, hangovers, skin-tags, grazes, etc. But this policy does not apply in *Death vs. Loss of Limbs*. Here, there is some number of people facing the prospect of losing their legs where we should save them at Bob's expense. Our judgments in these two cases are quite different. There is a class of views called *Limited Aggregation* that attempts to reconcile these judgments.<sup>2</sup> On these views, when one set of claims held by some individuals competes against a different set of claims held by another group of individuals, we *sometimes* satisfy the larger aggregate of claims. For example, we satisfy the larger aggregate of claims in *Death vs. Loss of Limbs*. But in *Death vs. Liver Spots*, we save Ann no matter how many people we can prevent from having liver spots. We thus resist aggregation in that case.

Limited Aggregation may be contrasted with *Pure Aggregation* and *Anti-Aggregation*. Pure Aggregation tells us always to aggregate and satisfy the larger aggregate of claims. In both *Death vs. Liver Spots* and *Death vs. Loss of Limbs*, there is some number of people facing the smaller affliction where we ought to help them rather than the person facing death. Anti-Aggregation, on the other hand, tells us that we should always act in favor of the individual with the strongest claim. Hence, we should save both Ann and Bob, no matter what. Both Pure Aggregation and Anti-Aggregation lead to counter-intuitive claims. Limited Aggregation is thus an attempt to occupy a middle ground that satisfies our moral judgments.

A central challenge of Limited Aggregation is explaining why we should aggregate in some cases, but not others. Attempts have been made by Alex Voorhoeve, David Lefkowitz, and Victor Tadros, among others.<sup>3</sup> For the most part, they have defended some *comparative* approach to the *Relevance View* (from here on, simply the Comparative Approach), which is roughly the conjunction of two ideas: firstly, claim A counts against claim B if and only if claim A is relevant to claim B (the Relevance View); secondly, claim A is relevant to claim B if and only if their relative strengths are close enough (the comparative aspect). However, it has been argued that the Comparative Approach has several counterintuitive implications such as requiring us to reject plausible principles about rational choice. In turn, some have defended the Comparative Approach against these implications.

In this paper, we offer an alternate version of Limited Aggregation, which is independently plausible and which avoids these criticisms. We call this version the Threshold Approach. The Threshold Approach is a *non-comparative* version of the Relevance View and states that there is a non-relative threshold that separates claims. We term claims that fall above the threshold *super-threshold* claims, and we term claims that fall below the threshold *sub-threshold* claims. Super-threshold claims are relevant to all other claims for the purposes of aggregation. Sub-threshold claims are not relevant to

<sup>2</sup>This term seems to have originated with Tomlin, "On Limited Aggregation." For some support for these judgments in public opinion, see Alex Voorhoeve, "Balancing Small against Large Burdens," *Behavioural Public Policy* 2, no. 1 (2018): pp. 125–42.

<sup>3</sup>Voorhoeve, "How Should We Aggregate Competing Claims?"; David Lefkowitz, "On the Concept of a Morally Relevant Harm," *Utilitas* 20, no. 4 (2008): pp. 409–23; Victor Tadros, "Localized Restricted Aggregation," in *Oxford Studies in Political Philosophy*, ed. David Sobel, Peter Vallentyne, and Steven Wall (Oxford: Oxford University Press, 2019), pp. 171–203. For an earlier statement of the Comparative Approach, see T. M. Scanlon, *What We Owe to Each Other* (Cambridge, MA: Belknap Press, 1998), pp. 238–40.

super-threshold claims. When some sub-threshold claims compete with other sub-threshold claims, the Threshold Approach has different options; our preference is for a version that holds that it is permissible to satisfy either set of sub-threshold claims. The Threshold Approach can in principle involve more than one threshold. In this paper, we propose a single threshold based on S. Matthew Liao's theory of Fundamental Conditions, which sets apart certain goods, capacities, and options as things to which human beings have a special claim.<sup>4</sup>

Some writers have considered the Threshold Approach, but they have summarily rejected it.<sup>5</sup> For instance, David Lefkowitz describes a principle along the following lines: "All of the harms within a given category are relevant to one another, and the relationship of transitivity holds between them. However, no harm in category 1 is relevant to any harm in category 2."<sup>6</sup> However, he then rejects it in the same paragraph, for reasons we shall discuss later. Our aim in this paper is to develop our version of the Threshold Approach and argue that it is much more plausible than one may initially think.

Here's how this paper will go. In part II, we outline the Comparative Approach and criticisms of it. In our view, those criticisms reveal devastating problems for proponents of the Comparative Approach. In part III, we outline the Threshold Approach and argue that it avoids many of the problems associated with the Comparative Approach. In part IV, we respond to some objections to the Threshold Approach.

Before we proceed further, it is worth making one clarificatory note. As we see it, Limited Aggregation is primarily about the deontic status of certain choices, not the axiological status of the results of those choices.<sup>7</sup> That is, our arguments here are concerned with what you ought to do, rather than which outcomes are better or worse. Just because we judge, for instance, that you ought to save Ann in *Death vs. Liver Spots*, does not mean that it is better for Ann to survive than for billions of people to have liver spots. We shall remain neutral on which outcomes are actually better or worse, from some impartial point of view, in most of the cases we discuss.

---

<sup>4</sup>S. Matthew Liao, "Human Rights as Fundamental Conditions for a Good Life," in *Philosophical Foundations of Human Rights*, ed. Rowan Cruft, S. Matthew Liao, and Massimo Renzo (Oxford: Oxford University Press, 2015), pp. 79–100.

<sup>5</sup>Not everyone rejects the Threshold Approach. Dale Dorsey, for example, offers an insightful defense of such a view ("Headaches, Lives and Value," *Utilitas* 21, no. 1 (2009): pp. 36–58). In addition, the Threshold Approach bears some similarities to other kinds of threshold views such as sufficientarianism and what Campbell Brown calls "threshold prioritarianism." See Roger Crisp, "Equality, Priority, and Compassion," *Ethics* 113, no. 4 (2003): pp. 745–63; Liam Shields, "The Prospects for Sufficientarianism," *Utilitas* 24, no. 1 (2012): pp. 101–17; Campbell Brown, "Priority or Sufficiency ... or Both?," *Economics and Philosophy* 21, no. 2 (2005): pp. 199–220. As we shall see, our project differs from these other views in a number of ways. For instance, many of these projects are concerned with well-being and are explicitly axiological, that is, they are interested in which outcome is better or worse (see also Shlomi Segall, "What Is the Point of Sufficiency?," *Journal of Applied Philosophy* 33, no. 1 (2016): pp. 36–52). In contrast, the Threshold Approach is concerned with certain fundamental conditions that human beings *qua* human beings need to pursue a good life rather than their well-being, and we are interested in duties and rights rather than which outcome is better or worse. Moreover, a key aim of our paper is to defend the structural advantages of a Threshold Approach over a Comparative Approach to Limited Aggregation, which these other projects do not address.

<sup>6</sup>Lefkowitz, "On the Concept of a Morally Relevant Harm," p. 413.

<sup>7</sup>See also Tomlin, "On Limited Aggregation."

## II. The Comparative Approach

As we have said, on the Comparative Approach, a claim is relevant to a stronger claim if and only if their relative strengths are close enough (the comparative aspect). For instance, Voorhoeve proposes that “A claim is relevant if and only if it is sufficiently strong *relative to the strongest competing claim*” (our italics).<sup>8</sup> The Comparative Approach faces some serious objections. The most salient ones involve two-stage problems, where we are presented with two or more groups of people whose claims you can satisfy. Claims are then either added to or subtracted from those groups, leading to counterintuitive conclusions that violate attractive principles of rational choice. For starters, Patrick Tomlin offers the following case:

### *Two-Stage Case 1:*

Stage 1: You can save Group A or Group B. Group A consists of one dying person.

Group B consists of ten people facing severe impairment.

Stage 2: Ten people facing mild impairment are added to both groups.<sup>9</sup>

Tomlin stipulates that mild impairment is relevant to severe impairment, and that one hundred claims against mild impairment are enough to outweigh ten claims against severe impairment. At the same time, ten claims against severe impairment outweigh one claim against death, while mild impairment is not relevant to death. Given this, at stage 1, you should save Group B, since ten claims against severe impairment outweigh one claim against death. However, when we add ten people facing mild impairment to both groups at stage 2, according to the Comparative Approach, the net strength of Group A’s claims increases. This is because the claims of those facing mild impairment added to Group A are relevant, as they compete with those facing severe impairment in Group B. However, the net strength of claims in Group B does not increase. This is because the claims of those facing mild impairment added to Group B are irrelevant, as they compete with a claim against death in Group A. Given this, the Comparative Approach would require us to save Group A after stage 2, despite the fact that equal claims were added to both groups. This is a very strange implication.<sup>10</sup>

Tomlin’s *Two-Stage Case 1* highlights a general problem with understanding the Comparative Approach. It violates what Korbinian Rürger calls Separability\*, which states that “If two distributions have subdistributions between which we are indifferent, then our choice between them depends only on what they are like in other subdistributions.”<sup>11</sup> To elaborate, let us say that we should be indifferent between the ten people facing mild impairment added to Group A and the ten others facing mild impairment

<sup>8</sup>Voorhoeve, “How Should We Aggregate Competing Claims?,” p. 66.

<sup>9</sup>Tomlin, “On Limited Aggregation,” p. 244.

<sup>10</sup>It is worth noting that Tomlin considers two different versions of the Comparative View, which he labels “Anchor by Strength” (where each claim is compared to the strongest claim in the entire case) and “Anchor by Competition” (where each claim is compared to the strongest claim it competes with). For brevity’s sake, we discuss here only Anchor by Competition. Tomlin also raises serious objections to Anchor by Strength (p. 244).

<sup>11</sup>Rürger, “Aggregation with Constraints,” p. 466. Rürger distinguishes between Separability and Separability\*. Separability states that “If two distributions have the same well-being level for one person, then our choice between them depends only on what they are like for other people.” The difference between the two is that Separability does not apply in *Two-Stage Case 1*, as there isn’t one person who is in both Group A and Group B.

added to Group B. According to Separability\*, the ten people facing mild impairment in both groups should “cancel out” each other at stage 2, leaving us in a position identical to where we were at stage 1. But the instructions provided by the Comparative View contradict this.

In addition to Tomlin, other writers have also presented related issues for the Comparative Approach. For instance, Norcross raises another two-stage problem, which we have adapted to be more similar to the other cases mentioned in this paper:

*Two-Stage Case 2:*

Stage 1: You have a choice between saving one life, one thousand people from severe impairment, or a million people from mild impairment.<sup>12</sup>

According to some versions of the Comparative Approach, you might be required to save the one thousand from severe impairment, assuming that mild impairment is irrelevant to the loss of a life, while saving a thousand people from severe impairment is relevant to the loss of a life. But now stage 2 kicks in:

Stage 2: On the way to saving those thousand people, it becomes impossible to save the one life – perhaps they die or a giant boulder gets in the way.<sup>13</sup>

Now the mild impairment becomes relevant, as it competes with severe impairment. Suppose that given a choice between saving a thousand people facing severe impairment, or a million people facing mild impairment, you should save the million. But why should the removal of one option make any difference, when you were not going to choose that option in the first place? Again, this is a very strange result. Indeed, it violates a common principle of rational choice known as Contraction Consistency: if some element of subset A of B is best in B, then it is best in A.<sup>14</sup> In Norcross’ case, within the set of saving one life, saving a thousand people from severe impairment, and saving a million from mild impairment, you ought to choose to save the thousand. Given this, Contraction Consistency means you should also save the thousand in any subset of those three options that contains this option. But this contradicts the result recommended by the Comparative Approach.

Norcross offers a further objection from Transitivity.<sup>15</sup> Imagine a choice between saving three groups. Group A consists of one person facing death. Group B consists of ten people facing severe impairment. Group C consists of a hundred people facing mild impairment. Assume again that mild impairment is not relevant to death while ten severe impairments outweigh one death. Given this, as Norcross points out, the Comparative Approach seems to imply that we should save Group B over Group A and Group C over Group B, but we should save Group A over Group C. Thus, we

<sup>12</sup>Alastair Norcross, “Two Dogmas of Deontology: Aggregation, Rights, and the Separateness of Persons,” *Social Philosophy and Policy* 26, no. 1 (2009): pp. 83–84.

<sup>13</sup>Ibid., pp. 83–84.

<sup>14</sup>Christopher J. Tyson, “Cognitive Constraints, Contraction Consistency, and the Satisficing Criterion,” *Journal of Economic Theory* 138, no. 1 (2008): p. 52. Contraction Consistency is also known as Independence of Irrelevant Alternatives in some contexts. Sen calls this principle  $\alpha$ . Amartya Sen, *Collective Choice and Social Welfare: An Expanded Edition* (Cambridge, MA: Harvard University Press, 2017), p. 63.

<sup>15</sup>Norcross, “Comparing Harms,” p. 139.

have a situation where  $A > C > B > A$ . This appears to violate the principle of Transitivity, which states that if  $B > A$  and  $C > B$ , then  $C > A$ .

Attempts have been made to salvage the Comparative Approach from these concerns.<sup>16</sup> One strategy involves denying the importance of Transitivity and Contraction Consistency with respect to deontic claims.<sup>17</sup> However, a theory that does not violate Transitivity and Contraction Consistency seems to have a dialectical edge over a theory that violates either, given the pretheoretic plausibility of these conditions. Another strategy involves rethinking the Comparative Approach.<sup>18</sup> For instance, Victor Tadros has defended what he calls *Local Relevance*, according to which even when a claim is irrelevant compared to a much stronger claim in a particular competition, it may still be relevant when weighed against yet another claim.<sup>19</sup> However, Joe Horton has pointed out difficulties with such a view.<sup>20</sup> As we shall shortly see, our non-comparative approach to Limited Aggregation offers a chance to escape these difficulties altogether.

### III. The Threshold Approach

The two-stage problems raised by Tomlin, problems with Transitivity, and other issues we have seen so far are symptomatic of accounts of Limited Aggregation that invoke comparative relevance. Most versions of Limited Aggregation ask us to consider whether some claim is strong enough, relative to some other claim. Owing to the comparative nature of such an approach, it is not surprising that a claim can be made relevant or irrelevant by the addition or removal of other claims. This easily leads to counterintuitive conclusions in multi-stage cases.

<sup>16</sup>One approach, which we will not discuss, involves arguing that decision theory principles, such as Separability\*, simply don't apply in the sorts of cases under consideration. For more, see Kirsten Mann, "The Relevance View: Defended and Extended," *Utilitas* 33, no. 1 (2021): pp. 101–10.

<sup>17</sup>For a discussion on this, see Tomlin (2017). For an author who rejects Transitivity in axiological cases, see Larry S. Temkin, *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning* (Oxford: Oxford University Press, 2011).

<sup>18</sup>Rüger defends a different version of the Comparative Approach, what he calls *Aggregation with Constraints*, which says that if one person has much more at stake than any person whose interests conflict with hers, we ought to do whatever serves the interests of that person. Otherwise we ought to do what realizes the best outcome. Interestingly, Rüger's account ends up violating his own Separability\*. Consider *Two-Stage Case 3*:

Stage 1: You have a choice between saving Group A, which consists of one person facing death, or Group B, which consists of a million people facing mild impairment.

Stage 2: Ten people facing severe impairment are added to Group A, while one person facing severe impairment is added to Group B.

Here, Aggregation with Constraints tells us to save Group A at stage 1, since the strongest individual claim (against death) is much stronger than the strongest claim with which it competes (mild impairment). But at stage 2, it seems that we are required to save Group B instead. This is because the addition of one person facing severe impairment in Group B means that the strongest claim in Group A (death) is no longer much stronger than the strongest claim with which it competes (severe impairment). At stage 2, Aggregation with Constraints entails that we should now "activate" the claims of the other people facing mild impairment and do what realizes the best outcome. This conclusion violates Separability\*, as the claim of the one person facing severe impairment added to Group B should be cancelled out by the claim of one of the ten people facing severe impairment in Group A.

<sup>19</sup>Tadros, "Localized Restricted Aggregation."

<sup>20</sup>Horton, "Always Aggregate."

The remedy that we propose is the Threshold Approach, which involves a non-comparative threshold that separates claims into super-threshold claims and sub-threshold claims. As mentioned earlier, super-threshold claims are relevant to all other super-threshold claims for the purposes of aggregation, while sub-threshold claims are not relevant to super-threshold claims. When some sub-threshold claims compete with other sub-threshold claims, as we shall shortly explain, our view is that it is permissible to satisfy either.

It might be asked, where is the threshold? There are a number of possible answers here, and the Threshold Approach is not committed to any specific theory. Further, our goal in this paper is not to defend any specific threshold or set of thresholds. Our goal is to show that the Threshold Approach avoids serious problems associated with the Comparative Approach to the Relevance View and is a plausible contender as a version of the Relevance View. Nevertheless, for illustrative purposes, we will use S. Matthew Liao's theory of Fundamental Conditions.<sup>21</sup> This theory picks out a list of goods, capacities, and options that are necessary for us, *qua* human beings, to pursue basic activities, which enjoy a special relationship with living a good life. These basic activities include pursuing "deep personal relationships with, for instance, one's partner, friends, parents, children; knowledge of, for example, the workings of the world, of oneself, of others; active pleasures such as creative work and play; and passive pleasures such as appreciating beauty."<sup>22</sup> On Liao's view, you do not need to pursue all the basic activities to have a good life, but it is impossible to live a good life without pursuing at least some of them. These activities are special because they are important to our lives as a whole *qua* human beings.<sup>23</sup> Basic activities may be contrasted with activities that are important to people *qua* individuals (such as pursuing an acting career), or activities that are only important to a small slice of a person's life (such as eating luxurious food).

A fundamental condition is a necessary ingredient, *qua* human beings, for the pursuit of some basic activity. Candidates include fundamental goods such as "food, water, and air"; fundamental capacities such as "the capacity to think, to be motivated by facts, to know, to choose an act freely (liberty), to appreciate the worth of something, to develop interpersonal relationships, and to have control of the direction of one's life (autonomy)"; and fundamental options, including "the option to have social interaction, to acquire further knowledge, to evaluate and appreciate things, and to determine the direction of one's life."<sup>24</sup> According to Liao, you have a right to all these fundamental conditions, even if you do not need to use them all to pursue a good life.<sup>25</sup> This is because when you deprive someone of a fundamental condition, or if you fail to provide it, they are unable to pursue some basic activity. This sets fundamental conditions apart from other goods that individuals may make claims to. Having a

<sup>21</sup>See Liao, "Human Rights as Fundamental Conditions for a Good Life" for a further articulation of the view and for a detailed defense.

<sup>22</sup>Ibid., p. 81.

<sup>23</sup>Ibid., p. 80.

<sup>24</sup>Ibid., p. 82.

<sup>25</sup>Ibid., p. 82. It may be helpful to note that Liao's theory is concerned with the goods, capacities, and options that human beings are entitled to. This theory is compatible with the idea that all humans are entitled to items within the set of fundamental conditions, even if they are unable to pursue a good life for other reasons. Hence, on this theory, a severely cognitively impaired person would still have a claim to any fundamental condition which others could reasonably provide even if it may be difficult for this person to use these fundamental conditions to pursue a good life.

liver spot, for instance, is slightly annoying, but it does not limit your ability to pursue basic activities. All other things being equal (excluding special responsibilities, etc.), I am required to ensure that you have access to the fundamental conditions, in a way that I may not be required to help you treat your liver spots.

Within the context of health, Liao has further distinguished between basic health issues and non-basic health issues.<sup>26</sup> Basic health issues are related to the organismic functions of a creature (e.g., respiration, digestion, absorption, metabolism, circulation, etc.) that are required for fundamental capacities, such as the capacities to think, sense, move about, and operate in the world around us. Non-basic health issues are those that do not pertain to the fundamental capacities. Examples include minor aesthetic issues such as liver spots, minor health problems such as a small cut or broken fingernail, and enhancements.

The distinction between basic and non-basic health issues helps us identify a plausible threshold for the relevance of claims. Among other things, basic health issues form the basis for super-threshold claims, while non-basic health issues sometimes only form the basis for sub-threshold claims. Here are some examples. Claims against death, paralysis, and the loss of limbs (and even parts of limbs, like fingers) fall above the threshold. Quite obviously, you cannot pursue basic activities if you are dead. Paralysis would also make it practically impossible to pursue at least some basic activities, such as many kinds of active pleasures.<sup>27</sup> Similarly, the loss of limbs, and parts of limbs, can affect your basic mobility such as your ability to manipulate objects and interact with them, which is important for pursuing many basic activities. A very bad case of the flu would also fall above the threshold, since influenza viruses can undermine your organismic functioning, which you need for thinking, moving about, and so on. Liver spots, grazes, broken fingernails, and hangovers, on the other hand, are non-basic health matters as they do not affect your fundamental capacities. A mild headache may fall above or below the threshold, depending on its etiology. For example, if a mild headache is symptomatic of a malignant brain tumor, it would impinge on basic health and thus fall above the threshold. But if a mild headache is the result of eating ice cream too quickly, it would not affect basic health and would fall below the threshold.

It is worth noting here that there could be claims that do not clearly constitute claims to fundamental conditions. For example, if a person suffers a cut, it could end up healing quickly because it is not that deep, but it could also become infected and eventually lead to the loss of some fundamental capacities because it is deep enough. In the next section, we shall point out that this raises the issue of vagueness or indeterminacy and discuss how the Threshold Approach can address it.

Besides health, Liao's fundamental conditions also include other capacities such as the capacity to think and the capacity to have interpersonal relationships. Here it may be helpful to point out that the basic/non-basic distinction can apply to these other capacities as well. For instance, consider a person's capacity to think. On the Fundamental Conditions theory, some aspects of thinking such as means-end reasoning, planning, solving day-to-day problems, interpreting the actions of others, and

<sup>26</sup>S. Matthew Liao, "Health (Care) and Human Rights: A Fundamental Conditions Approach," *Theoretical Medicine and Bioethics* 37, no. 4 (2016): pp. 259–74.

<sup>27</sup>This does not mean that a disabled person cannot live a good life since a disabled person often can pursue some basic activities. Also, it is worth noting that the Fundamental Conditions theory is about what human beings are entitled to and does not assess whether a specific person's life is better or worse because she cannot pursue some basic activities.

justifying our behavior would be considered basic since they are essential for pursuing the basic activities. At the same time, other aspects of thinking such as being able to do calculus or quantum mechanics would not be needed for pursuing the basic activities and would therefore be non-basic. Or, to use another example: the capacity for interpersonal relationships. Once again, there is a qualitative difference between “being able to make some friends” and “being able to become the mayor of a city.” Here the social skills required for making some friends is basic, while the charisma required to be a politician is not basic.

Even if you are not sure about the Fundamental Conditions approach that underlies our threshold, other theories in political philosophy do converge on a similar threshold. To give an example, Martha Nussbaum specifies a list of capabilities related to life, bodily health, bodily integrity, senses, imagination, thought, emotions, practical reason, affiliation, other species, play, and control over one’s environment.<sup>28</sup> It may be possible to set the threshold at whatever conditions are required to maintain such capabilities. Such a theory would agree that claims to life, against threats of significant bodily harm, against threats of disability, and to important goods such as food and water fall above the threshold. It would also agree that claims against liver spots, or claims to recreational goods like ice cream or a World Cup subscription fall below the threshold. This general inter-theory agreement raises both our confidence that there is a threshold that divides claims and our confidence in approximately where the threshold lies.

To see how the Threshold Approach is a theory of Limited Aggregation, consider again the examples presented at the outset:

*Death vs. Liver Spots:* We can either save Ann from death or prevent some number of different people from having liver spots for the rest of their lives.

*Death vs. Loss of Limbs:* We can either save Bob from death or prevent some number of different people from losing both their legs.

In *Death vs. Liver Spots*, Ann’s claim is a super-threshold claim, while liver-spot claims are sub-threshold claims. On the Threshold Approach, sub-threshold claims are not relevant to super-threshold claims. Therefore, this approach says that we should save Ann.

In *Death vs. Loss of Limbs*, Bob’s claim to be saved from death and the other people’s claims to be saved from losing both of their legs are all super-threshold claims since they all pertain to basic health matters. Hence, the Threshold Approach says that there is some number of people facing the prospect of losing their legs where they should be saved at Bob’s expense. As noted earlier, satisfying this pair of judgments makes the Threshold Approach a theory of Limited Aggregation.

Next, let us turn our attention to a major advantage that the Threshold Approach has over the Comparative Approach. In particular, the Threshold Approach provides straightforward solutions to the problems with Separability\*, Contraction Consistency, and Transitivity raised by Tomlin’s and Norcross’ two-stage problems. Primarily, this is because the relevance of a claim under the Threshold Approach depends on the *internal* properties of the claim, that is, the strength of a claim depends on the properties of the claim itself, e.g., whether it affects basic health. This may be

<sup>28</sup>Martha C. Nussbaum, *Creating Capabilities: The Human Development Approach* (Cambridge: Belknap Press, 2011).

contrasted with the Comparative Approach, where relevance is determined by some *external* property, that is, the relative strength of a claim depends on how it fares when compared to another claim.

We can demonstrate this advantage of the Threshold Approach by looking at the kind of cases presented by Tomlin and Norcross. But before proceeding, a quick word on our presentation below. As we have seen, in the literature, it is common to employ examples that involve more abstract negative experiences such as pain and impairment. These abstract negative experiences are then taken to compete against more concrete ones such as death. We are concerned that using abstract experiences can misleadingly give the impression that negative health states such as minor discomforts and severe pain or minor impairment and severe impairment differ *only* in degrees. However, we believe that negative health states can in some cases differ *qualitatively* and not just in degrees. In addition, we believe the moral disvalue of pain should not be abstracted from the conditions that cause them. Pain is not bad just because of how it feels. It is also bad because it may be the effect of an organism's suffering from a potentially harmful condition, which itself is bad (as we shall see later, we believe that a similar point can be made regarding the duration of a harmful condition). Indeed, some types of pain, like a hangnail, are non-basic health matters, while other types of pain, like a broken bone, involve basic health matters. Elsewhere, Justin Klocksiem makes a slightly different point to the same effect. According to him, some negative experiences are mere "discomforts," which are intrinsically bad, but "in a trivial and insignificant way," while other negative experiences involve genuine pain.<sup>29</sup> For this reason, we will avoid examples that involve abstract negative experiences such as pain and impairment. We suspect that most people would regard our point to be terminological and would not quibble with us. But we also realize that there may be other people who believe that negative health states differ necessarily only in degrees. Shortly, we shall respond to objections that are motivated along such lines.

Let us consider the following adaptation of *Two-Stage Case 1*:

Stage 1: You can save Group A or Group B. Group A consists of one dying person.

Group B consists of ten people who will lose both of their legs.

Stage 2: A hundred people facing liver spots are added to both groups.

Let us assume that claims against death and against losing both legs fall above the threshold, while claims against liver spots fall below the threshold. Let us also assume that ten claims against losing both legs are collectively stronger than one claim against death. The Threshold Approach says you should save Group B at stage 1. At stage 2, the Threshold Approach maintains this judgment. A claim against liver spots falls below the threshold. So, the people added at stage 2 do not make a difference here.

Hence, unlike the Comparative Approach, the Threshold Approach, owing to its non-comparative nature, does not violate Separability\*. Recall that "If two distributions have subdistributions between which we are indifferent, then our choice between them depends only on what they are like in other subdistributions."<sup>30</sup> The Threshold Approach satisfies this condition, because the relevance of a person's claim does not depend on what it is compared to. If some person's claim is below the threshold,

<sup>29</sup>Justin Klocksiem, "How to Accept the Transitivity of Better Than," *Philosophical Studies* 173, no. 5 (2016): pp. 1309–34. See also Dorsey, "Headaches, Lives and Value" for a similar point.

<sup>30</sup>Rüger, "Aggregation with Constraints," p. 466.

then her counterpart in a different group will also have a sub-threshold claim. For the purposes of choosing between those two groups, the sub-threshold claims will always be irrelevant.

Recall Norcross's worry that the Comparative Approach appears to violate Contraction Consistency. To see how the Threshold Approach does not violate this principle, consider an adapted version of Norcross' thought experiment (*Two Stage Case 2*):

- Stage 1: You have a choice between saving one life, one thousand people from losing both of their legs, or a ginormous number of people from liver spots.  
 Stage 2: The choice to save the one life becomes unavailable.

Suppose that at stage 1, you should save the thousand people from losing both of their legs. The Threshold Approach implies that you should maintain that decision at stage 2. This is because, at both stages, the claims against liver spots fall below the threshold and should not be counted. Hence, unlike the Comparative Approach, the Threshold Approach does not violate Contraction Consistency. The reason again is that, on the Threshold Approach, the relevance of claims is not comparative. If a claim ought to be satisfied within a larger set, it ought to be satisfied within any subset of that larger set.

Likewise, the Threshold Approach also avoids problems with Transitivity (assuming they exist). Suppose that Group A consists of one person facing death, Group B consists of ten people who will lose both of their legs, and Group C consists of a hundred people with some sub-threshold claim such as liver spots. The Threshold Approach would say that you ought to save Group B over Group A, Group A over Group C, and Group B over Group C. Thus,  $B > A > C$ . This does not violate Transitivity.

Given that the Threshold Approach does not violate Separability\*, Transitivity, or Contraction Consistency, this should give it a dialectical edge over the Comparative Approach, which has been shown to violate these principles of rational choice.

Finally, let us consider how the Threshold Approach would treat sub-threshold claims that compete with one another. As we shall see, the Threshold Approach has different options. The one we prefer says that sub-threshold claims are claims that we are permitted, but not required, to satisfy. Call this the Sub-threshold Claims Are Not Requiring (SCANR) version of the Threshold Approach. The rationale for SCANR is as follows: Sub-threshold claims are not very weighty moral claims. Indeed, in the case of health, for example, they involve non-basic health matters such as liver spots, grazes, broken fingernails, and so on. It is true that helping someone to not to have these conditions is beneficial for that person. But liver spots, grazes, broken fingernails, and so on, do not affect a person's basic ability to pursue a good life. For this reason, sub-threshold claims are not very weighty moral claims. Given this, it seems that it is permissible to satisfy them, but one should not be required to satisfy them. In other words, we can distinguish between "deontic" reasons and "supererogatory" reasons.<sup>31</sup> Deontic reasons are reasons that obligate and require us to perform certain acts, other things being equal. Supererogatory reasons are reasons that permit us to perform certain good or beneficial acts, but these reasons do not obligate or require us to perform these acts, other things being equal. Since sub-threshold claims are morally not

<sup>31</sup>Supererogatory reasons can refer to reasons that are not very weighty or reasons that would be too demanding to satisfy, given the circumstances. Here we have in the mind the former.

very weighty, SCANR says that they give us supererogatory reasons. Super-threshold claims, on the other hand, are morally very weighty because they affect a person's basic ability to pursue a good life. As such, super-threshold claims give us deontic reasons, other things being equal.

To illustrate SCANR, consider a simple case in which you can easily wave a wand and prevent one person from having liver spots for the rest of her life. According to SCANR, since this is a sub-threshold claim, one is permitted but not required to satisfy this claim. Similarly, consider a case in which some sub-threshold claims compete with other sub-threshold claims:

X: Prevent five people from having liver spots for the rest of their lives.

Y: Prevent ten people from having liver spots for the rest of their lives.

Again, since liver spots are sub-threshold claims, SCANR also implies that it is permissible to satisfy either X or Y. Again, this is because having liver spots will not prevent someone from being able to pursue a good life.

Some philosophers, including proponents of the Comparative View, are likely to disagree on this point. In particular, they are likely to say that we can have deontic reasons to satisfy even relatively weak claims such as claims against liver spots. This brings us to the other version of the Threshold Approach, Sub-threshold Claims are Requiring (SCAR), which says that absent countervailing considerations, we are required to satisfy sub-threshold claims. For instance, in a case in which you can easily wave a wand and prevent one person from having liver spots for the rest of her life, SCAR says that you are deontically required to help in this case because there is no countervailing reason not to help. Likewise, in a case where you can prevent five people from having liver spots (X) or ten people from having liver spots (Y), SCAR says that you are deontically required to help the ten people, since the claims at issue, i.e., liver spots, are of the same strength. In other words,  $Y > X$ .

The issue for SCAR though is when stronger claims are added to the mix. For instance, suppose that the possibility of being able to save an additional life is added to both X and Y.

X\*: Save A's life and prevent five people from having liver spots for the rest of their lives.

Y\*: Save B's life and prevent ten people from having liver spots for the rest of their lives.

In this sort of case, SCAR faces a dilemma. Either it proposes a course of action which seems unfair and disrespectful to A, or it violates Separability\*.

Let us consider the first horn of the dilemma. A proponent of SCAR might say "When only sub-threshold claims compete with each other, satisfy the group with the highest aggregate of sub-threshold claims. If both groups have equal super-threshold claims, still satisfy the group with the largest aggregate of sub-threshold claims." Thus, in the above case, they might say that you are required to perform action Y\*. However, if one says that we should give B's life priority in virtue of the fact that we can prevent more people from having liver spots in Y\*, one might think that this is unfair and disrespectful to A, given that liver spots are weak, sub-threshold claims. In fact, some defenders of a Comparative Approach have made a similar point. For instance, Frances Kamm asks us to consider a case in which you can save Albert's life in one

group or save Betty's life and cure Collin's sore throat in another group. Kamm argues that if Collin did not have a sore throat, it would be permissible to save either Albert or Betty. According to Kamm, the presence of Collin's sore throat should not make a normative difference in this case because sore throats in this case are what she calls "irrelevant utilities," which are utilities that are not relevant when a much stronger claim is present.<sup>32</sup> Since sore throats should be irrelevant in this case, Kamm argues that it would be unfair and disrespectful to Albert if Albert loses all chances of being saved just because of the presence of Collin's sore throat. In other words, Collin's sore throat should not be used as a tie-breaker in such a case. Hence, there are good reasons to believe that it would be unfair to give B's life priority in virtue of the fact that we can prevent more people from having liver spots in  $Y^*$ . If so, this is an issue for a proponent of SCAR who holds that  $Y^* > X^*$ .

Suppose instead a proponent of SCAR agrees that it is permissible to satisfy either  $X^*$  or  $Y^*$ , in other words,  $X^* \sim Y^*$ . The second horn of the dilemma would kick in. In particular, their view would now violate Separability\*. The reason is that SCAR initially says that  $Y > X$ . Then when a life is added to each group, it now says  $X^* \sim Y^*$ , which is a violation of Separability\*. In contrast, SCANR does not violate Separability\*. This is because with respect to  $X$  and  $Y$ , SCANR says that  $X \sim Y$ , deontically speaking. And with respect to  $X^*$  and  $Y^*$ , SCANR says that  $X^* \sim Y^*$ , deontically speaking. So Separability\* is not violated. For this reason, we prefer SCANR, even though we recognize that it might be counterintuitive to some that we have no deontic reasons to satisfy sub-threshold claims. Nevertheless, proponents of the Threshold Approach have multiple options available to them, and the Threshold Approach is plausible as long as one of these options is acceptable.

#### IV. Potential concerns

Now for some worries with the Threshold Approach. One initial concern pertains to the examples we have used in our defense of the Threshold Approach, in particular the use of liver spots. We have said that claims against liver spots are not relevant to claims against more serious impairments, like the loss of limbs. But you may point out that liver spots are so mild that even those who defend the Comparative Approach would not consider them relevant to the loss of limbs. Thus, you may worry that we have stacked the deck by choosing cases which are obviously and trivially true.<sup>33</sup>

There are a couple of reasons for our choice of examples. First, we have opted against the use of cases like "mild impairment," because as explained earlier, we think they are too abstract. Second, we have chosen liver spots precisely because they are so obviously trivial that most would agree that no number of liver spots is relevant to the loss of a limb. Importantly, proponents of the Comparative Approach would agree with this. Given this, we do not think that this use of examples unfairly favors our view.

Another potential concern is that a single super-threshold claim, even if it is quite small, can outweigh any number of sub-threshold claims even when the benefits are significant. For instance, suppose that we can give a million people a million dollars each, but they all have the full set of fundamental conditions already and so they would only

<sup>32</sup>Frances Kamm, *Morality, Mortality Vol. I: Death and Whom to Save from It* (New York: Oxford University Press, 1993), p. 146. See also Bastian Steuwer, "Aggregation, Balancing, and Respect for the Claims of Individuals," *Utilitas* 33, no. 1 (2021): pp. 17–34, for such a view.

<sup>33</sup>Credit to Christian Barry for raising this worry.

use the money on luxury goods;<sup>34</sup> or, we can save a person from losing a limb. The Threshold Approach seems to imply that we should save the person from losing a limb. While this conclusion may strike some people as radical, we find this conclusion acceptable. As we noted in the introduction, we are concerned with the deontic status of our choices. We are interested in what we are required to do, not what is best for us to do. And it seems correct to us that you are not morally required to give people millions of dollars if they will only use it to buy luxury goods.

A further concern is that the Threshold Approach is committed to an arbitrary threshold. Lefkowitz makes this kind of argument. As he writes, “The obvious problem with this conception of morally relevant harms is that it will inevitably be arbitrary where exactly along the scale of continuous harms we establish a boundary between different categories of harm.”<sup>35</sup> There are two ways to understand the arbitrariness concern. The first is that there is no principled way to draw a threshold or, as Lefkowitz puts it, a boundary between different categories of harms. Lefkowitz suggests this objection, writing, “Why is it, given that there is so little difference between L and M, that L is a category 1 harm, while M is a category 2 harm?”<sup>36</sup> We believe that the response to this concern is already built into our explication of the Threshold Approach based on the Fundamental Conditions Theory. We have provided a substantive account for why there is a boundary between different categories of harms. A person’s basic capacity to pursue a good life is a morally salient notion – it is something that can give us reasons to act in specific ways. It is therefore plausible that our reasons to respond to claims regarding fundamental conditions are different and stronger than our reasons to respond to claims to non-fundamental conditions. So there is a principled distinction between different categories of harms.

Another interpretation of the arbitrariness concern is as a concern from vagueness. Perhaps there is no way to specify a *specific* threshold without appealing to guesswork or some vague and shaky intuitions. For example, imagine a series of increasingly serious cuts, each of increasing length but equal depth. Do we draw the threshold between a 0.9-inch cut and a 1-inch cut? Or perhaps between a 1-inch cut and a 1.1-inch cut? Or perhaps between a 0.95-inch cut and a 0.96-inch cut? There is some arbitrariness in the choice of a specific threshold. Similarly, even with the Fundamental Conditions Theory, some vagueness persists. For instance, it is not entirely clear how much mobility a person needs for her to pursue basic activities. Now, vagueness can be a practical or epistemic issue. That is, the fact that some threshold is vague can make it very difficult to identify it, or make it very difficult for agents to achieve unanimity on a threshold, and justify any particular threshold. However, our project here is primarily conceptual. We are not claiming that thresholds within the Threshold Approach will be easy to identify. We are claiming that somewhere, in the space of claims, there is a threshold.

But the objection from vagueness could be that because of vagueness, we should doubt that there is a threshold. There are two responses to this objection. Firstly, the problem of vagueness is not unique to the Threshold Approach. As a matter of fact, all Relevance Views, including Comparative Approaches, face it. For example, on a Comparative Approach, is a 10-second headache relevant to a lifelong headache? If not, what about an 11-second headache, and so on? Presumably, at some point, a

<sup>34</sup>We thank an anonymous reviewer for this thought.

<sup>35</sup>Lefkowitz, “On the Concept of a Morally Relevant Harm,” p. 413.

<sup>36</sup>*Ibid.*, p. 413.

Comparative Approach would say that an N-second headache will become relevant to a lifelong headache. If so, the problem of vagueness would also arise because it is vague when an N-second headache becomes relevant to a lifelong headache. Indeed, the problem of vagueness seems to exist beyond the topic of Limited Aggregation. For instance, most non-consequentialists believe that it is impermissible to kill one person in order to save two, three, or even ten other people. But at some point, it is permissible to kill in order to save a much larger number of lives.<sup>37</sup> It might be thought that it is vague where the threshold for permissible killing lies. In any case, since the problem of vagueness afflicts the Comparative Approach as well, we should not hold this objection specifically against the Threshold Approach. As mentioned previously, it is not our goal to defend a specific threshold, but to point out that the Threshold Approach has important advantages over the Comparative Approach.

Secondly, just because the line between two concepts is vague does not mean that those two concepts do not exist. For example, imagine a light slowly changing from orange to red. At some point, the light is orange, but at another, the light is red. But it is vague when this change occurs. This does not mean that we have made some categorical error; that “orange” and “red” are illusory or useless concepts. It just means that the world contains vague concepts and blurry lines between them – a fact that most of us live with every day. Similarly, the line between a super-threshold claim and a sub-threshold claim might be a vague one. Occasionally, we may have to make arbitrary distinctions between them for the purposes of guiding our decisions. But this does not mean that the principles behind those distinctions are flawed.

You may press the objection from vagueness further by pointing out that vagueness is a more serious problem for the Threshold Approach than the Comparative Approach because the Threshold Approach assumes – indeed, requires – a discontinuity within the space of claims. For any threshold, there is a sub-threshold claim that just barely falls below the threshold, such that it is irrelevant to a super-threshold claim that falls just barely above the threshold. It is this discontinuous nature that enables the Threshold Approach to satisfy various principles of rational choice in a number of cases. However, this discontinuous nature carries with it certain implications. In particular, it implies that a claim just above the threshold should be satisfied instead of many claims just below the threshold. Here’s a specific example. Suppose that on the Fundamental Conditions theory, a claim against a 1-inch cut barely falls above the threshold, because it still affects some basic health matter. At the same time, suppose that a claim against a 0.9-inch cut falls below the threshold, but just barely, because it is just superficial enough not to affect basic health. The Threshold Approach implies that you should save one person from a 1-inch cut over many people from a 0.9-inch cut. This judgment may seem troubling. We can call this the *adjacent claims* concern.

We believe that there’s a principled reason for why you should not be too troubled by this concern. First, some claims that seem adjacent might not actually be that close. As we have argued, the threshold marks a division between basic health issues and non-basic health issues. Basic health issues pertain to the organismic functioning of an individual, which in turn affects whether that individual is able to pursue basic activities. This makes basic health issues a fundamental condition, and non-basic health issues

<sup>37</sup>This is the view of Threshold Deontology, which holds that standard deontological constraints do not apply when the consequences of upholding those constraints are sufficiently great. For more discussion of threshold deontology, see Lawrence Alexander, “Deontology at the Threshold,” *San Diego Law Review* 37 (2000): pp. 893–912.

a non-fundamental one. If a 1-inch cut falls above the threshold, then even if it does so barely, this still means that a 1-inch cut is the kind of thing that can affect your basic organismic functions. Perhaps it affects some important muscles which in turn affects your basic mobility. Perhaps it causes pain so severe that you struggle to think about much else. Or perhaps a 1-inch cut is deep enough to create a sufficiently high chance of further complications such as a life-threatening infection. No matter how we fill in the details, given that it is above the threshold, a 1-inch cut is something that can affect your basic organismic functions. For similar reasons, the Threshold Approach also implies that there is some number of people facing 1-inch cuts such that we should save them over one person facing death (assuming that these claims are equally urgent). In our view, this implication is also acceptable. The reason again is that we are supposing that a 1-inch cut is the kind of thing that can affect your basic organismic functions and therefore basic health.

At the same time, if a 0.9-inch cut falls below the threshold, then even if it does so barely, this means that a 0.9-inch cut cannot affect your basic organismic functions, however unpleasant it may be. Once we fill in the details, we see that there is a stark moral difference between a claim against a 0.9-inch cut and a claim against a 1-inch cut. A claim against a 0.9-inch cut is a claim against an unpleasant experience. A claim against a 1-inch cut, on the other hand, is a claim to basic organismic functions and therefore basic health. These claims differ not just in degree of strength, but in the kinds of claims they are. In other words, some things that look like adjacent claims are not in fact that close to each other. Given this, it seems correct that a claim just above the threshold should be satisfied instead of many claims just below the threshold, other things being equal.

Some people may deny that thresholds between claims can really be so sharp, such that a claim against a 1-inch cut falls above it, while a claim against a 0.9-inch cut falls below it. There are two responses available to us here. The first is to maintain that there really are such sharp thresholds, although it is very hard in practice to determine where they are. It seems plausible to say that for any agent, there is a point at which some injury or ailment is so bad that they cannot exercise their fundamental capacities and pursue basic activities. For instance, suppose that being able to move about unaided is a fundamental capacity. For any agent, there is plausibly some threshold of injury or ailment that will prevent her from being able to move about unaided. While it might be hard to find this threshold, and while such thresholds can differ slightly from person to person, they exist nonetheless. The second response is that even if no sharp threshold exists, the Threshold Approach is compatible with a “fuzzy” threshold. For example, suppose that there is a fuzzy threshold from 0.5- to 1.5-inch cuts, such that a claim against a cut longer than 1.5-inches is above the threshold, while a claim against a cut shorter than 0.5-inches is below the threshold. On this picture, claims against cuts between 0.5- and 1.5-inches are indeterminate, and it is not clear whether they are relevant to claims above the threshold. While this indeterminacy is not very satisfying, it is nevertheless compatible with the Threshold Approach. Even if claims against cuts between 0.5- and 1.5-inches are indeterminate, the Threshold Approach can still make distinctions between claims against cuts below 0.5-inches and cuts above 1.5-inches.

An adamant skeptic may argue that the problem persists. For any goods, capacities or options that are on Liao’s fundamental conditions list, a potential loss with respect to those fundamental conditions could also lie anywhere on a spectrum from extremely small to extremely large. Thus, there must be two truly adjacent claims somewhere

along the spectrum which sit on opposite sides of the spectrum. For example, take the issue of mobility. There may be a continuous spectrum of possible effects on mobility ranging from the trivial to the devastating. Even if the Threshold Approach gives a plausible list of goods, capacities, and options, it may not help specify a non-arbitrary threshold with respect to the magnitudes of losses of those goods, capacities, and options such that losses above the threshold can be aggregated and losses below it cannot.

As far as we can tell, this is a variant of Lefkowitz's arbitrariness concern, this time applied to the fundamental conditions such as basic mobility. As with our response to the arbitrariness concern, we can either interpret this concern as one concerning vagueness, or as the claim that we have no principled method of drawing a threshold in the spectrum of goods, capacities, and options. In response to the first interpretation, the vagueness concern, our response is again that just because the line between two concepts may be vague does not mean that those two concepts do not exist. In response to the second interpretation, we believe that it is already built into the Fundamental Conditions theory to draw a line separating goods, capacities, and options that are basic and fundamental from those that are not basic and fundamental. In the case of mobility, for example, we would argue that there is a qualitative difference between non-basic mobility and basic mobility. Non-basic mobility pertains to actions that are not needed, *qua* human beings, in order to pursue the basic activities. Some examples might include being able to ski or dunk a basketball. In contrast, basic mobility involves actions that are necessary, *qua* human beings, in order to pursue the basic activities. It seems clear that if a person could not move at all, that person's ability to pursue the basic activities would be seriously hampered. While there is certainly a spectrum of mobility on which both "being able to move at all" and "being able to dunk a basketball" exist, there is clearly a qualitative difference between them. Somewhere between these two activities, there is a threshold, whether sharp or fuzzy, between basic mobility and non-basic mobility. Wherever this threshold is, it is plausible that claims grounded in non-basic mobility are irrelevant to claims grounded in basic mobility.

At this point, someone might accept that there could be qualitative differences between a 1-inch cut and a 0.9-inch cut, but they may claim that the argument does not extend to the case of pain.<sup>38</sup> In particular, it might be argued that the Threshold Approach implies that a claim against  $N$  amount of pain can be below the threshold, while a claim against  $N+1$  amount of pain can lie above it. However, so this argument goes, this seems implausible since  $N+1$  pain is just a bit worse than  $N$  pain. In response, as noted earlier, we believe that the moral disvalue of pain should not be abstracted from the conditions that cause it and that there can be qualitative differences between seemingly similar negative experiences. Once we recognize this, it is not implausible to think that  $N$  pain could be qualitatively different from  $N+1$  pain.

Thus far, we have stated the challenge to the Threshold Approach in terms of the types of the goods one may be deprived of. However, the challenge can also be stated in terms of the duration for which one is deprived of certain goods. It might be thought that the Threshold Approach will have difficulty handling cases involving duration. Someone who is concerned with this aspect of the Threshold Approach might offer the following cases:

<sup>38</sup>We thank Patrick Tomlin for this suggestion.

*Paralysis for Life vs. Paralysis for One Minute:* We can either save Jane from being paralyzed for life or prevent some number of different people from being paralyzed for one minute.

*Paralysis for Life vs. Paralysis for Five Years:* We can either save Vincent from being paralyzed for life or prevent some number of different people from being paralyzed for five years.

In *Paralysis for Life vs. Paralysis for One Minute*, it might be thought that we should save Jane regardless of how many people we can potentially save from one minute of paralysis. In *Paralysis for Life vs. Paralysis for Five Years*, however, it seems that there is some number of people facing the prospect of being paralyzed for five years where we should save them at Vincent's expense. If this is correct, this pair of judgments would be similar to other types of Limited Aggregation judgments. However, it might be thought that, on the Threshold Approach, paralysis would be a super-threshold claim, since paralysis prevents people from being able to pursue a good life. If so, it seems that the Threshold Approach would have to say that there is some number of people facing the prospect of being paralyzed for one minute where we should save them instead of Jane. If so, the Threshold Approach would seem to reach the wrong verdict regarding *Paralysis for Life vs. Paralysis for One Minute*.

To see why this example does not undermine the Threshold Approach, it is helpful to question its underlying assumption that being saved from paralysis is necessarily a super-threshold claim. This assumption is problematic because there are different kinds of paralysis with different causes and, arguably, only some types of paralysis give rise to super-threshold claims. For instance, hereditary spastic paraplegia, cerebral palsy, and spina bifida would all give rise to super-threshold claims since they are congenital and can impair fundamental capacities. Paralysis after an injury or a stroke would also give rise to super-threshold claims since they are the result of interference with organismic functioning. Likewise, primary periodic paralysis, which is a group of rare diseases that temporarily make muscles stiff, weak, or unable to move, would also give rise to a super-threshold claim, even if it lasts for just one minute, because it is also congenital. On the other hand, a one-time, minutes-long episode of sleep paralysis caused by sleep deprivation would not give rise to a super-threshold claim since occasional sleep deprivation does not affect basic health. Given that different types of paralysis have different etiologies, either the minute of paralysis is the result of something that affects basic health or it is not. If it is the result of something that affects basic health, e.g., if it is a case of primary periodic paralysis, then even if it lasts for just one minute, a case can be made that there is some number of people facing the prospect of having primary periodic paralysis where we should save them instead of Jane. On the other hand, if the minute of paralysis is the result of something that does not affect basic health, such as temporary sleep paralysis, then the Threshold Approach can also reach the verdict that we should save Jane regardless of the number of people we may be able to save from temporary sleep paralysis.

Some people may assert that it is not just the underlying causes of a condition that matter; sometimes the symptoms themselves matter. For instance, all else equal, paralyzes themselves are equally important. Even so, this does not challenge the Threshold Approach. Suppose that two cases of paralysis are identical in all respects to each other except for their duration. The Threshold Approach has the resources to group them into different kinds of claims. This is because the duration of a condition can

affect whether the condition prevents a sufferer from pursuing basic activities. For example, a one-time episode of paralysis that lasts a minute is unlikely to affect one's ability to pursue basic activities. But a longer episode of paralysis that lasts hours, or a chronic case of short recurring episodes, can force a person to sit out a number of important basic activities. If so, the Threshold Approach can state that paralysis for one minute is not a super-threshold claim, while longer episodes are super-threshold claims. More generally, for any given condition, there may be a duration  $T$  such that a case lasting  $T$  or longer provides the sufferer with a super-threshold claim to assistance.

At this point, other people might attempt to put forward a version of the adjacent-claims concern, this time with respect to duration.<sup>39</sup> In particular, suppose that paralysis of  $T$  duration ought to be addressed because it is a super-threshold claim. Some people might worry that this could imply that we should prevent one paralysis of  $T$  duration over a very large number of the same paralysis but of  $T-1$  duration. Our response to this concern is similar to what we have said in response to the adjacent-claims concern. For one thing, if we are correct that duration can also affect one's ability to pursue the basic activities, then there is a fact of the matter whether  $T-1$  duration affects one's ability to pursue the basic activities or not. If it does not, then the claims are not as close as we may have thought and there are principled reasons why we should give priority to one paralysis of  $T$  duration. Alternatively, one might be skeptical that the threshold to the claims involving duration could be so sharp. Suppose that this is the case. As before, the Threshold Approach can accommodate a "fuzzy" threshold, such that the threshold lies over a range of durations for a given condition. For example, paralysis of  $T$  duration is a super-threshold claim, while paralysis of  $T-50$  is a sub-threshold claim. Everything in between is indeterminate. In this scenario, claims above the threshold lie some distance away from claims below the threshold. If so, the Threshold Approach need not commit us to saving one case of paralysis lasting  $T$  over many similar cases lasting  $T-1$ . If all of this is right, the Threshold Approach can reach the correct verdict even in sets of cases that differ only in terms of duration.

## V. Conclusion

Let us look once more at the (now expanded) list of the options available to us.

1. Pure Aggregation: When groups of claims compete, we have reasons to satisfy the greatest aggregate of claims.
2. Anti-Aggregation: When groups of claims compete, we have reasons to satisfy the group with the largest single claim.
3. Limited Aggregation, which includes:
  - A. The Comparative Approach, where a claim is relevant to a stronger claim if and only if their relative strengths are close enough.
  - B. The Threshold Approach, where we are required to satisfy the largest aggregate of super-threshold claims, and we are permitted but not required to satisfy the sub-threshold claims.

Each one of these options has costs associated with it. Pure Aggregation and Anti-Aggregation both yield some very counterintuitive conclusions when tested

<sup>39</sup>We thank Patrick Tomlin for this point.

against either *Death vs. Liver Spots* or *Death vs. Loss of Limbs*. The Comparative Approach avoids these problems. However, it yields other counterintuitive results when faced with a variety of multiple-stage cases. These counterintuitive results are symptomatic of underlying ways in which the Comparative Approach violates plausible principles such as Separability\*, Contraction Consistency, and Transitivity. The violation of these principles does not necessarily make a theory false, but it seems that these violations are costs for a theory that violates them, at least on the face of it.

Our own proposal, the Threshold Approach, does not violate these plausible principles, owing to its non-comparative nature. Importantly, the Threshold Approach is also explanatory; by setting a threshold using the Fundamental Conditions theory, there are principled reasons for supposing that there is a substantive difference between different kinds of claims. For these reasons, we believe that the Threshold Approach is at least a strong competitor, if not the most attractive choice, among the options on the table.

**Acknowledgements.** We would like to thank Alex Voorhoeve, Patrick Tomlin, Christian Barry, Ralf Bader, Tom Parr, Aidan Penn, Patrick Wu, Adam Lerner, Bastian Steuwer, Michael Rabenberg, Nicholas Tilmes, Chris Lernpaß, audiences at University of Freiburg, the University Center for Human Values at Princeton University, the Centre for Moral, Social, and Political Theory at the Australian National University, and the Centre for Ethics, Law and Public Affairs at University of Warwick, and two anonymous reviewers for their very helpful comments on earlier versions of this paper.