Ethic Theory Moral Prac (2016) 19:71–82
DOI 10.1007/s10677-015-9662-9

The Dimensions of Consequentialism: Reply to Schmidt, Brown, Howard-Snyder, Crisp, Andric and Tanyi, and Gertken

Martin Peterson [1]

Accepted: 31 October 2015 / Published online: 23 November 2015

Abstract In this article I respond to comments and objections raised in the special issue on my book *The Dimensions of Consequentialism*. I defend my multi-dimensional consequentialist theory against a range of challenges articulated by Thomas Schmidt, Campbell Brown, Frances Howard-Snyder, Roger Crisp, Vuko Andric and Attila Tanyi, and Jan Gertken. My aim is to show that multi-dimensional consequentialism is, at least, a coherent and intuitively plausible alternative to one-dimensional theories such as utilitarianism, prioritarianism, and mainstream accounts of egalitarianism. I am very grateful to all contributors for reading my book so closely and for devoting time and intellectual energy to thinking about the pros and cons of multi-dimensional consequentialism.

Keywords Multi-dimensional consequentialism · Utilitarianism · Prioritarianism · Equality · The dimensions of consequentialism

I feel immensely honoured by the contributions to this special issue on *The Dimensions of Consequentialism*. I sincerely thank all authors for reading my work so closely and for devoting time and intellectual energy to thinking about the pros and cons of multi-dimensional consequentialism. In what follows I will respond to what I take to be the main points of each article. However, before I do so, I would like to give an example of how multi-dimensional consequentialism differs from other versions of consequentialism. I shall return to this example on several occasions in my discussion.

✉ Martin Peterson
martinpeterson@tamu.edu

[1] Department of Philosophy, Texas A&M University, 4237 TAMU, YMCA Building 301, College Station, Texas 77843, USA

1 An Example

By definition, multi-dimensional consequentialists believe that an act's rightness or wrongness depends on several irreducible moral aspects, which typically belong to different dimensions.[1] For instance, an act that produces at least as much wellbeing as every alternative might be right with respect to one aspect (wellbeing), while an act that produces a perfectly equal distribution might be right with respect to another (equality). If no alternative is optimal with respect to all moral aspects, then multi-dimensional consequentialists believe that no alternative is, all things considered, entirely right. Moral rightness and wrongness come in degrees. If two or more moral aspects clash and no alternative is optimal (or suboptimal) with respect to all of them, then each alternative is *literally speaking* somewhat right and somewhat wrong.

In the example sketched in Table 1, alternative A maximises wellbeing. Utilitarians, therefore, believe that A is right and that B and C are wrong.[2] Strict egalitarians, if they exist, believe that only C is right while prioritarians would typically claim that B is the only right alternative.[3] Multi-dimensional consequentialists who believe that wellbeing and equality are two irreducible moral aspects will point out that while A is optimal with respect to total wellbeing it scores poorly with respect to equality. They also observe that the opposite is true of C, and while B is not optimal with respect to any of the two aspects, it scores pretty well with respect to both of them. All things considered, none of the three alternatives is therefore entirely right. All three alternatives are somewhat right and somewhat wrong. Perhaps B is almost entirely right (right to a degree close to 1 on a scale from 0 to 1) while A and C are, roughly speaking, half right and half wrong (right to degree 0.5).[4]

2 Moral Conflicts

By a moral conflict, Thomas Schmidt means a situation in which no act is optimal with respect to all moral aspects. As correctly pointed out by Schmidt, multi-dimensional consequentialists analyse such conflicts as a kind of moral dilemma: no matter what you do, you will do something that is at least a little bit wrong.

Schmidt argues that a better account of moral conflicts can be obtained by retaining the traditional binary notion of rightness, but *extend* the binary analysis by explicitly mentioning the conflicting moral aspects in the final analysis. For instance, Schmidt's proposal entails that in the example sketched in Table 1, each of the three alternatives is either right or wrong in the binary sense (perhaps B is right, and A and C wrong). Moreover, because there is a conflict between two moral aspects, equality and wellbeing, this clash deserves to be mentioned in the final analysis. Schmidt notes that Ross' notion of prima facie duties and Dancy's conception of contributory reasons can be interpreted as variations of this proposal.

[1] I distinguish between 'aspects' and 'dimensions'. A dimension is the conceptual space in which an aspect can vary. For instance, the area of a circle depends on only one aspect (its radius) whereas the area of a triangle depends on two aspects (its base and height). All three aspects are elements of the same dimension (length), but the length and mass of a tennis ball are different aspects which are also elements of different dimensions.

[2] This example is similar to the one I give in Chapter 1.

[3] Naturally, the prioritarian verdict will depend on the shape of the (concave) priority function.

[4] The technical aspects of multi-dimensional consequentialism can be developed in a number of ways. I discuss three proposals in the appendix to *The Dimensions of Consequentialism*.

Table 1 An example

Act	Alice's wellbeing	Bob's wellbeing
A	99	1
B	50	49
C	0	0

What Schmidt and I disagree about is what counts as a *final analysis* of a moral problem. To be more precise, our disagreement is about what one should expect to find on the very last line in an account of a moral conflict. Schmidt is quite liberal. I am considerably more conservative. Schmidt is willing to include a broad range of moral judgements in his final analysis, such as claims about normatively relevant factors as well as extensive assessments of their relative weight.

Contrary to Schmidt, I believe the final analysis of a moral conflict should be restricted to verdicts about each act's deontic status, that is, its rightness or wrongness. As emphasised above, such deontic verdicts may sometimes be non-binary. We should, perhaps, conclude that B in Table 1 is right to degree 0.9 (on a scale from 0 to 1) while A and C are right to degree 0.5.[5] I believe an advantage of my approach is that it preserves the *normative power* of 'right' and 'wrong' in a way that Schmidt's proposal fails to capture. To see this, imagine that Schmidt were to conclude that alternative B in Table 1 is right (although he would, of course, add that B is suboptimal with respect to both equality and wellbeing, and so on) while A and C are wrong. Now imagine that we replace B by a new alternative B', as in Table 2.

Alternative B' is clearly optimal with respect to all relevant aspects, so in his final analysis Schmidt would have to say that B' is right in the binary sense and add that there is no conflict between wellbeing and equality. However, what strikes me as odd is that the *moral difference* between B and B' is entirely accounted for by these supplementary judgments in Schmidt's analysis. The conclusion about the rightness and wrongness of B and B' is the same in both cases: both acts are right in the binary sense.

Why is this problematic? To start with, we *intuitively feel* that B' is better than B from a moral point of view, but this difference is not reflected in our deontic verdicts. The difference becomes visible only in our non-deontic, supplementary moral judgements.

It is also appropriate to consider an analogy with beliefs. If you accept a binary account of belief ('full belief' versus 'full disbelief') and are confronted with two situations in which you fully believe that it will not rain today, although you in the second situation have some reason for thinking the opposite, then one might expect this epistemic difference to lead to some difference in behaviour. One possibility could be to claim in the second case that you should bring an umbrella if you go for a walk. However, if you think more carefully about it, you will see that this is a bit odd. Why should someone who *fully believes* that it will not rain bring an umbrella when going for a walk? The agent's behaviour seems to indicate that she believes, at least to some degree, that it will rain.

In a similar vein, Schmidt's extended analysis of moral conflicts seems to entail that once it has been established that an act is right, the supplementary judgements will make no difference to the agent's behaviour. In fact, Schmidt admits that once

[5] See the appendix to *The Dimensions of Consequentialism* for a discussion of this numerical representation of degrees of rightness.

Table 2 An entirely right act

Act	Alice's wellbeing	Bob's wellbeing
B'	100	100

considered, his supplementary judgements about conflicting moral aspects 'do not have to be considered again when it comes to deciding how to act in the particular situation'. (Schmidt, this issue) The upshot is that you should *not* bring an umbrella if you fully believe that it will not rain, even if you have some good reason for thinking the opposite. There is, therefore, no practical difference between the two cases. We should expect the agent to behave in exactly the same way in both cases. This strikes me as implausible. If some moral aspects clash in one case, but not in another, we should expect there to be some practical difference between the two cases.

To this Schmidt could perhaps respond as follows: Imagine that we replace B' by a new alternative B'' in which Alice and Bob get 50 units each. From a multi-dimensional point of view, alternative B'' would also be entirely right (because it is optimal with respect to both aspects, wellbeing and equality). But B'' is clearly not as good as B', meaning that there is a moral difference between B' and B''. If multi-dimensional consequentialists insist that all moral differences should be accounted for in the final verdict, then it seems impossible to do so if B' and B'' are entirely right, but B' is better than B''.

My reply to this objection is that there *is* a moral difference between B and B', but not between B' and B'' It is, of course, true that B' brings about more wellbeing than B'', but B' and B'' are not alternatives to each other. We are comparing alternatives across different situations. The reason why B in the original situation is not right in the binary sense is that it is not optimal with respect to all applicable moral aspects. However, both B' and B'' have all the properties that make right acts entirely right: they are optimal with respect to every applicable aspect. What matters is that both B' and B'' are optimal with respect to both wellbeing and equality. Therefore, we do not have to, and should not, make any moral distinction between B' and B''.

3 Can We Have the Cake and Eat It Too?

Campbell Brown argues that there is a predicate, *right*, which is binary, and a comparative relation, *more right than*, which comes in degrees. According to Brown, both these notions of rightness are required for fully characterising an act's deontic status. Moreover, unlike the multi-dimensional account of degrees of rightness I defend, Brown's account is one-dimensional. He explicitly claims that the distinction between multi- and one-dimensional consequentialism is irrelevant to his thesis, but I think this is false.[6] I admit that what Brown claims about one-dimensional consequentialism might be true, but it would be a mistake to think that Brown's point applies to multi-dimensional consequentialism.

[6] Brown uses the term value pluralism. The distinction between value pluralism and multi-dimensional consequentialism is discussed in Chapter 1 of *The Dimensions of Consequentialism*.

The examples considered by Brown are all of the following type: You are faced with three alternatives, which will bring about either 1, 2 or 3 units of some good. The best option, to bring about 3 units, is according to Brown right in the binary sense and the two others are wrong in the same sense. It is, of course, *better* to bring about 2 units rather than 1, so to account for this difference Brown thinks we should extend the binary judgements with non-binary ones: it is *more right* to bring about 2 units than 1, although both options are *less right* than bringing about 3 units.

As far as I can see, the structural features of Brown's proposal are very similar to what we would end up with if we were to interpret Mill literally. Mill famously claimed that 'actions are right in proportion as they tend to promote happiness, wrong as they tend to produce the reverse of happiness.'[7] So if you bring about 2 units of happiness but could have brought about 3, you do something that is right to degree 2/3, and so on. This literal interpretation of Mill entails a non-binary (and one-dimensional) account of rightness.[8] However, in all other parts of *Utilitarianism*, Mill operates with a traditional binary notion of rightness. So it seems that a strict Millian would, perhaps, accept both parts of Brown's position.

For the moment, I shall leave it open whether Brown's one-dimensional view is normatively plausible. As I see it, the key issue is whether Brown's position would also be acceptable to multi-dimensional consequentialists. Contrary to what Brown himself claims, I think the answer is no.[9] If an act's deontic status depends on several irreducible moral aspects, then the binary part of Brown's view should be rejected. To see this, it is helpful to consider how multi-dimensional theorists are supposed to make decisions in light of their all-things-considered moral verdicts. Brown agrees with Thomas Schmidt that the best options are always right in the binary sense. Therefore, one should always perform one of those options.

However, if the multi-dimensional theorist were to claim that the morally best alternatives are always right in the binary sense, then the multi-dimensional theorist would *also* have to conclude that one of the best alternatives should always be performed. For the multi-dimensional theorist, but perhaps not for the one-dimensional one, this is an unfortunate conclusion. When faced with a clash between conflicting moral aspects, as the one illustrated in Table 1, multi-dimensional consequentialists have strong reason to believe that a rational and morally motivated decision maker would distribute her acts over all alternatives that are right to *some* degree, i.e. randomise over all alternatives that are right to some non-zero degree. By doing so, the agent gives all conflicting aspects their due. Randomisation reflects the complexity of the moral landscape faced by the multi-dimensional theorist. If the agent were to conclude that some aspects are totally outweighed by others and conclude that the best options are right in the binary sense and that one of them should, therefore, be performed with probability one, then her behaviour would not reflect the fact that each moral aspect corresponds to a verdictive moral reason.[10]

Brown objects to this that it is, in a strict sense, *impossible* to perform the type of randomised acts I believe multi-dimensional consequentialists should perform in order to give

[7] Mill (1861/1969: 210).
[8] Note that this literal interpretation of Mill presupposes a ratio measure of happiness.
[9] See the first paragraph of Section 2 in his article.
[10] See Chapter 6 of *The Dimensions of Consequentialism*.

all aspects their due. His point is that 'randomisation is not distribution'.[11] His argument for this goes as follows:

> To see this, consider again belief. Suppose my evidence is neutral between P and ¬P. What should I do? One thing I could do—the *rational* thing, I would say—is distribute my credence equally, giving 1/2 to P and 1/2 to ¬P. Another thing I could do is toss a coin, giving 1 to P if it lands heads (and 0 to ¬P), or 0 to P if it lands tails (and 1 to ¬P). Clearly, the latter approach is not a way of partly believing each of two propositions. Whichever way the coin lands, I will believe one proposition completely, and the other not at all. But the latter approach is Peterson's randomisation. The former approach—the rational one—has no analogue in the case of action. (Brown, this issue)

Contrary to Brown, I believe there is no fundamental difference between distributing one's belief over several propositions and distributing one's action over several options. To be more precise, I think Brown misdescribes (or at least oversimplifies) the randomisation procedure I have in mind. If one were to first toss a coin and then perform either act P or ¬P, depending on the outcome, then this would resemble the case in which one tosses a coin and then fully decides to believe either P or ¬P. As Brown himself notes, this is not what it means to distribute one's belief over several propositions.

The following example might help to bring more light to the notion of randomisation I have in mind: Imagine that there are two canteens on campus, A and B, in which faculty and staff can have lunch. Each canteen is optimal with respect to some aspect, but not with respect to all. So no matter which canteen you go to, you are bound to go to one that is suboptimal with respect to at least one aspect. The *collective act* performed by all faculty and staff can be said to be distributed over both alternatives, meaning that each day about half of the group has lunch in each canteen (although the members of each group, of course, varies over time). Moreover, each agent's individual action over time will also be distributed over both alternatives. If you, for instance, have lunch on campus about 200 days per year you will go to each canteen about 100 times per year. But what about your individual action on a single occasion? Clearly, you cannot visit both canteens 'to a degree' on a single day. You either have lunch in canteen A or B, even if you take the decision after having tossed a coin.

My response to this objection is that we should think of the mental process that triggers your behaviour (say, the decision to have lunch in canteen B) as being part of the randomised act. Imagine that you shortly before you head out of your office to grab some lunch are asked to estimate your subjective probability that you will have lunch in canteen B. After considering your indeterminate preference between the two canteens, you conclude that your subjective probability that you will have lunch in canteen B is ½. Moreover, this subjective probability of yours will remain the same until you have actually entered one of the canteens; at that time it will, of course, be either 0 or 1. To randomise one's action and distribute it over several alternatives on a single occasion means that the agent is *uncertain* about which alternative she will perform in the near future. An agent who is in a deliberative mode and remains uncertain of what the outcome of this deliberate process will be until it is over is, on the account I propose, distributing her individual action on a single occasion over several alternatives.[12]

[11] Brown, this issue.

[12] Wolfgang Spohn (1977) and Isaac Levi (1989) have pointed out that it can be problematic to ascribe subjective probabilities to one's own choices. I address their worries in my (2006) and (2008).

4 One-Dimensional Views

Frances Howard-Snyder makes two points. Her first point is that rightness can, in her opinion, vary in degrees even if an act's rightness or wrongness depends on a single moral aspect. The assumption about multiple dimensions or aspects is therefore not essential for my claim about degrees of rightness. Her second point is that it is a mistake to think that an act cannot be entirely right if it is wrong with respect to at least one aspect. Even acts that are suboptimal with respect to some aspect can be entirely right.

I agree with Howard-Snyder that it would be conceptually possible to defend a non-binary account of rightness without presupposing a multi-dimensional view. As mentioned above, John Stuart Mill is an example of a one-dimensional theorist who, if interpreted literally, would be committed to a non-binary account of rightness. According to Howard-Snyder, a non-binary account could be advantageous for one-dimensional consequentialists because it would enable them to give more useful advice about moral matters. If I am, for example, convinced that you will not actually perform your best option, then a non-binary account would enable me to advise you to at least do your second best option, because I could say that while this option is not *entirely* right it is at least right to a *higher degree* than some other options.

I find all this quite convincing, but I believe Howard-Snyder and I disagree on what is ultimately at stake in this discussion. The non-binary notion of rightness Howard-Snyder proposes is clearly different from that advocated by multi-dimensional consequentialists. As Howard-Snyder sees things, there will always be at least one option that is entirely right. This is, however, something that is explicitly denied by multi-dimensional consequentialists. Multi-dimensional consequentialists believe that rightness varies in degrees because conflicting dimensions or aspects sometimes clash. In such cases, no option will be entirely right, as illustrated in Table 1. Because Howard-Snyder's non-binary notion of rightness differs in this fundamental respect from the multi-dimensional one, it seems that we could accept her point without revising any beliefs about what multi-dimensional consequentialists should say or believe about degrees of rightness.

Howard-Snyder's second point is that it is always entirely right to do one's best, even if the morally best alternative is suboptimal with respect to some applicable aspect. In her view,

> If you say that someone's action is partially wrong, you owe the agent or the listener some account of what the agent should have done. For example, a professor should be able to tell the student who earned a B^- how she could have earned an A. (Howard-Snyder, this issue)

Contrary to Howard-Snyder, I believe that genuine moral dilemmas are possible, meaning that doing one's best may *not* always be entirely right. For instance, if wellbeing and equality are separate aspects, and no alternative is optimal with respect to both (see Table 1), then this is a situation in which the best alternative is not entirely right. Instead of offering the agent an account of what the agent *should have done* in this situation, it is more appropriate to offer the agent an explanation of *why* the agent is facing a dilemma in which no alternative is entirely right.

I believe the same applies *mutatis mutandis* to student grades. Although nearly all students surely have the capacity to earn an A, it seems likely that not all do. When advising such students it would be misleading to try to convince someone who earned a B⁻ that he or she could have earned an A. Given the student's knowledge and mental capacity, it might very well hold true that there was no way in which the student would have been able to do what was required for an A. In this type of case, it seems that all the professor could be asked to do is to clearly state the criteria a student (act) has to meet in order to earn an A (be entirely right), but there is no need to claim that it would be possible for the student (act) to meet those criteria.

To steer clear of this type of objection, Howard-Snyder appeals to the influential idea that 'ought' implies 'can'. If you cannot perform an act that is optimal with respect to all aspects, then it is not the case that you ought to perform that alternative. Therefore, if we take 'ought' to mean 'right', this spells trouble for the multi-dimensional theorist who maintains that an act is (entirely) right if and only if it is optimal with respect to all aspects, no matter whether any such alternative is available. My response to this is that defenders of moral dilemmas have little reason to believe that 'ought' implies 'can'. So for the multi-dimensional consequentialist who believes that he or she can offer a consequentialist analysis of situations that involve moral dilemmas, it is far from clear that 'ought' implies 'can'. In fact, it seems that multi-dimensional consequentialists have given us a reason for thinking that 'ought' does not always imply 'can'.

Whether moral dilemmas are possible or not is, of course, an open question. However, I believe that the analogy with the student who cannot get an A, in the relevant sense of 'can', supports the intuition that dilemmas are possible. We have no reason to let the actual abilities of a student determine what we require for an A. Our standard should be objective and the same for all students. And this also seems to apply to acts. We have no reason to let the features of the acts open to us in a particular situation determine what is required of an act for it to be morally right. The standard for moral rightness is objective and the same across all situations. An act is entirely right if and only if it is optimal with respect to all aspects.

5 Does It Matter How One Arrives at a Moral Verdict?

Roger Crisp and I agree that John Stuart Mill's concept or rightness is best interpreted as a binary notion. The fact that a literal interpretation of a single sentence in *Utilitarianism* would entail a non-binary account does not change this. It merely shows that Mill chose his words badly on at least one occasion.

Crisp also argues that my rejection of the binary account is premature because, in his opinion, the standard account of rightness has all the resources required for accounting for the considerations I believe support the non-binary account. The clash between equality and wellbeing in Table 1, again, illustrates this point. As a multi-dimensional consequentialist I believe the moral complexity in Table 1 can only be accurately described by a non-binary notion of rightness. The traditional view that clashes between conflicting aspects can always be reduced to binary moral verdicts is too heavy-handed. Moral nuances that ought to be reflected in our final moral verdicts will invariably be omitted when a complex moral analysis is described by a binary moral judgement.

Crisp responds to this by claiming that one can account for the moral complexities triggered by clashes between conflicting moral aspects by distinguishing between the way an agent *arrives* at her moral verdict and the information that needs to be recognised *within the moral*

verdict itself. As he puts it, the multi-dimensional consequentialist's non-binary account of degrees of rightness,

> involves blindness to how the proponent of the standard view arrives at her verdict. She fully takes into account the significance of the failure to promote equality, but this is while she is *arriving* at her verdict. What we need from Peterson is an argument for recognizing *within some verdict itself* considerations that are relevant to it. Since proponents of the standard view recognize these considerations in arriving at their verdicts and needn't be taken to have forgotten them once they issue their verdict, I think such an argument would be hard to make. (Crisp, this issue)

Crisp's suggestion is somewhat similar to Schmidt's. They both believe that we can account for the complexity triggered by conflicting moral aspects by supplementing the traditional binary analysis with some additional information. Schmidt argues that the supplementary information should be a set of binary moral judgements that *describe the clash* between different aspects. Crisp's proposal is that we should focus on the *process* that lead up to the binary verdict.

Let us suppose, for the sake of the argument, that Crisp is right: it is indeed possible to account for all the moral nuances that multi-dimensional consequentialists consider to be important by mentioning *how* the agent arrived at her final (binary) verdict. An unwelcome implication of any theory of this type would be that rational and morally motivated decision makers will always be required to perform an act that is right in the binary sense. This is because some of the conflicting aspects will always 'win' in the end, and only those aspects will be reflected in our decisions. In the example with the two canteens discussed in my reply to Brown, I argued that each canteen (each somewhat right alternative) should be chosen with some non-zero probability. For similar reasons, multi-dimensional consequentialists who accept Crisp's point will not be able to make clashes between conflicting aspects visible in their choice behaviour. So the argument for why we should recognize '*within some verdict itself* considerations that are relevant to it', is that this enable us to justify a reasonable theory of decision making for multi-dimensional consequentialists.[13]

Crisp also makes several interesting comments about the value of equality. His worry is that 'egalitarianism is non-welfarist' and 'once you start allowing in non-welfarist values, there seems to be no obvious place to stop'.[14] In my discussion of equality in chapter four of *The Dimensions of Consequentialism*, I argue that unfair inequality can be bad for someone even if the person benefits from an unequal distribution. Whether this is a welfarist or non-welfarist account depends on how we define welfarism (I do not use this term myself), but I do think my account makes it clear where we should stop: If something is not good or bad *for* a person, then it does not matter from a moral point of view. All moral values, including equality, are individual values that make things go better or worse *for* a person.

[13] Op. cit.
[14] Op. cit.

6 Risk

Vuko Andric and Attila Tanyi discuss a series of alleged drawbacks of the account of risk offered in chapter five of *The Dimensions of Consequentialism*. Very briefly put, I argue that risk should be treated as an independent moral aspect on par with wellbeing and equality. While I admit that non-binary accounts of risk could be spelled out in numerous ways, I also develop my own preferred account, which is based on actualist intuitions (or, put in the terminology preferred by Andric and Tanyi, 'objectivist' intuitions). According to my account, the degree to which an act is right, with respect to risk, is equal to the probability with which it will bring about the greatest amount of wellbeing. My argument for this view is that I believe it tallies better with our considered moral intuitions than alternative views proposed in the literature, such as objective consequentialism, according to which an act is right if and only if it actually brings about the greatest amount of wellbeing, and subjective consequentialism, according to which an act is right if and only if no alternative would bring about a greater expected amount of wellbeing.

Andric and Tanyi question the claim that my multidimensional account tallies better with our considered moral intuitions than objective consequentialism (OC) and subjective consequentialism (SC). They present a list of five problems, some of which spell trouble for objective theories and some of which cause problems for subjective ones, and claim that 'all problems of OC and SC are problems for [the multidimensional account] as well'.[15]

I agree with nearly all points made by Andric and Tanyi.[16] Despite this, I do not think we have to conclude that there is something wrong with the multi-dimensional account. I am aware that this might sound a bit odd, so let me explain.

The reason why I think Andric and Tanyi's argument has less drastic implications than they seem to think is that they ignore some considerations that support the multi-dimensional account. When deciding what to believe about a philosophical issue, we should not merely consider arguments that speak against various views. Arguments that support each relevant view should also be examined. As far as I can see, Andric and Tanyi do not discuss any of the arguments supporting the views they discuss.

It is also worth keeping in mind that some of the objections put forward by Andric and Tanyi are formulated as conditional claims. For example, they argue that '*if* the Jackson case [one of their five problems] poses a problem for OC, then Peterson's proposal shares this problem.'[17] It is not clear to me what we should make of this claim if we think the Jackson case poses *no* problem for OC.

So far I have discussed the general structure of Andric and Tanyi's argument. However, let me also say a few words about what I believe to be the most important of the five problems they bring up, viz. the multi-dimensional analysis of the Jackson case. (See Table 3 in their article.) I agree it might be somewhat counterintuitive to say that it is entirely wrong to opt for the partial cure. Despite this, I think the multi-dimensional conclusion that the two other cures are half-right and half-wrong are more attractive than the traditional objectivist conclusion that one of the two is right in the binary sense although we cannot know which. So the multi-dimensional view is at least an improvement over the traditional objectivist view *as seen from the objectivist's point of view*. I do not believe that this particular form of multi-dimensional account will score better than all other views with respect to all relevant considerations, but I do believe it scores better in the Jackson case than the traditional objectivist account.

[15] Andric and Tanyi, this issue
[16] The only exception is their analysis of the Jackson case, as I explain below.
[17] Andric and Tanyi, this issue.

7 Is It Bad to Receive More than One's Fair Share?

Jan Gertken discusses my multi-dimensional analysis of Taurek's lifeboat problem. Multi-dimensional consequentialists favour a *mixed* solution to this type of problem. If you can rescue either *m* or *n* people, and there are no morally relevant differences among the groups, it is sometimes–but not always–entirely right to toss a fair coin. However, if the difference in size between the two groups is large enough, then no alternative will be entirely right. So no matter what you do, you will do something that is somewhat wrong, meaning that many lifeboat problems can be construed as non-binary moral dilemmas in which at least two moral aspects clash: fairness and the wellbeing of each individual life at stake.

A crucial premise of my argument for the mixed analysis is that unfairness is bad not just for people who receive *less* than they are entitled to (that is, are rescued with an unfairly low probability), but also for those who receive *more* than their fair share (are rescued with an unfairly high probability). It is this premise of my argument that Gertken seeks to challenge in the first half of the paper. However, he does not give any explicit argument for rejecting it. If I understand him correctly, he believes that we have a strong and robust intuition that unfairness is bad only for people who receive *less* than they are entitled to. If this is his argument, I disagree with him.

In chapter four of the book, I discuss an analogy with the Olympic Games. It is, I argue, better to be defeated in the Olympics in a 'clean' competition in which no illegals steroids are used because then there will be no unfair inequalities among the athletes. As far as I can understand, Gertken is willing to accept this part of the argument. If you lose because others get an unfair advantage, then that is bad for you. Now, my point is that it also seems plausible to say that it is better to *win* the Olympics if you do so without using any illegal steroids, compared to the case in which you win an unfair victory. That is, your victory is worth more *for you* if it is a fair victory. Note that this is not a psychological claim about how you would perceive your victory. It is a claim about the value of the victory itself.

I take it that intuitions about unfairness in sports support the idea that unfairness can be bad for someone even if that person benefits from the unfair distribution of some good. This is just a claim about the value of the unfairness itself. It is fully possible, and even likely, that it is better for most of us to win a gold medal in an unfair competition than to finish in the middle of the pack in a fair completion.

Is this analogy relevant for the general ethical claim I am trying to defend? I am willing to admit that my analogy is just a single data point, but I think it at least shows that the intuition that unfairness is bad only for people who receive less than they are entitled to is not as strong and robust as Gertken seems to believe. I think this indicates that Gertken's criticism needs to be supported with some additional argument.

In the final sections of his paper, Gertken sketches an interesting alternative account of the mixed solution, which he claims to be superior to mine. The main difference is that he allows for interpersonal aggregation of reasons: The claims each of us have to be rescued can be interpersonally aggregated, but because fairness also matter one should not always opt for the utilitarian solution and save the greatest number. For instance, in a choice between saving 1000 or 1001, the total amount of fairness gained by tossing a fair coin matters more than the potential loss of one extra life.

Gertken's preferred version of the mixed solution strikes me as a very reasonable account of what a multi-dimensional consequentialist who believes in interpersonal aggregation should think, but as I explained earlier I do not believe that this interpersonal route is feasible. Despite this, I think Gertken's proposal is a valuable contribution to the multi-dimensional discussion over fairness.

References

Levi, I. (1989) 'Rationality, Prediction, and Autonomous Choice', *Can J Philo* 19 (Suppl), 339–362. Reprinted in Levi, I. (1997): *The Covenant of Reason*, Cambridge University Press.

Mill JS (1861/1969) *Collected works of John Stuart Mill*. University of Toronto Press, Toronto.

Peterson M (2006) 'Indeterminate preferences', *Philos Stud* 130(2):297–320.

Peterson M (2008) *Non-bayesian decision theory*. Springer.

Peterson M (2013) *The dimensions of consequentialism*. Cambridge University Press.

Spohn W (1977) 'Where Luce and Krantz do really generalize Savage's decision model', *Erkenntnis* 11:113–134.