

Economics and Philosophy

<http://journals.cambridge.org/EAP>


Additional services for ***Economics and Philosophy***:

Email alerts: [Click here](#)

Subscriptions: [Click here](#)

Commercial reprints: [Click here](#)

Terms of use : [Click here](#)



VAULTING INTUITION: TEMKIN'S CRITIQUE OF TRANSITIVITY

Alex Voorhoeve

Economics and Philosophy / Volume 29 / Issue 03 / November 2013, pp 409 - 423
DOI: 10.1017/S0266267113000321, Published online: 15 October 2013

Link to this article: http://journals.cambridge.org/abstract_S0266267113000321

How to cite this article:

Alex Voorhoeve (2013). VAULTING INTUITION: TEMKIN'S CRITIQUE OF TRANSITIVITY. *Economics and Philosophy*, 29, pp 409-423 doi:10.1017/S0266267113000321

Request Permissions : [Click here](#)

VAULTING INTUITION: TEMKIN'S CRITIQUE OF TRANSITIVITY

ALEX VOORHOEVE

London School of Economics, UK
a.e.voorhoeve@lse.ac.uk

Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning, Larry S. Temkin. Oxford University Press, 2012, xxi + 613 pages.

How to rank distributions of benefits and harms? In this book, Larry Temkin addresses this question in detail. Its core claims are two. First, the goodness of a distribution is sometimes 'essentially comparative' – it sometimes depends on which alternative distribution(s) it is compared to. Second, there are many cases in which our intuitions are at odds with the transitivity of 'all things considered better than' and these cases give us reason to doubt that this relation is transitive. (Transitivity holds that if some alternative a_3 is better than a_2 , and a_2 is better than a_1 , then a_3 is better than a_1 .)

It is of great importance whether these claims are true. First, in social choice theory, it is often assumed that the goodness of an alternative is assessed solely on the basis of its associated well-being outcomes (see, for example, Fleurbaey *et al.* 2009: 259). If Temkin is right, then this assumption is unwarranted.

Second, transitivity plays a key role in allowing our choices to be guided by our notions of value. Transitivity is a necessary condition for assigning a precise numerical value to each alternative. It also guarantees that there is at least one alternative that is not strictly worse than some other alternative. It therefore ensures that our choices can be consistent with our notion of betterness, in the following sense: we can avoid

I am grateful to Larry Temkin for discussion of his ideas and to Jason Alexander, Ken Binmore, Richard Bradley, Francesco Guala and Michael Otsuka for comments. I also thank Princeton University's Center for Human Values for the Faculty Fellowship during which I wrote this critical notice.

choosing an alternative that we regard as strictly worse than some other alternative. Moreover, when transitivity is violated because we have cyclical betterness judgements, we cannot ensure such consistency. For example, if we hold that a_3 is better than a_2 , and a_2 is better than a_1 , but a_1 is better than a_3 , then we regard every alternative as worse than some other alternative, so that our choices cannot be guided by our notion of betterness.

The book draws on material from Temkin's papers challenging transitivity that span more than two decades. These papers kicked off a lively discussion.¹ Part of the value of Temkin's book is that it draws together, substantially revises, and elaborates on this material and offers his responses to critics. What is also significant is the way in which the book connects its critique of transitivity to the idea that the goodness of a distribution is essentially comparative. In a more comprehensive manner than he has done before, Temkin argues that the fact that the goodness of a distribution is essentially comparative explains when and why transitivity fails.

Temkin's book is rich and long. In what follows, I shall focus on two of his central cases. I shall argue that the first of these offers some support for the idea that the goodness of a distribution may depend on what other distributions are feasible. But I shall also argue that neither case threatens transitivity. Contrary to Temkin, I shall therefore argue that transitivity holds even when the goodness of a distribution depends on the feasible set of which it is a member.

1. MANY-PERSON SPECTRUM CASES

Chapter 1 provides an overview of the book. Chapter 2 kicks off the arguments with a discussion of interpersonal trade-off cases. Temkin (pp. 30–32) posits two principles for deciding such cases:

The First Standard View – Trade-offs between Quality and Number are Sometimes Desirable: In general, an outcome where a larger number of people have a lower quality benefit is better than an outcome where a smaller number of people have a higher quality benefit, *if* the number receiving the lower quality benefit is 'sufficiently' greater than the number receiving the higher quality benefit and *if* the differences in the initial situations of the people and the degrees to which they are benefited are not 'too' great.

The Second Standard View – Trade-offs between Quality and Number are Sometimes Undesirable even when Vast Numbers are at Stake: If the quality of

¹ See, among others, Temkin (1987; 1996), Norcross (1997, 1999), Rachels (1998), Binmore and Voorhoeve (2003), Broome (2004: ch. 4), Qizilbash (2005), Carlson (2005) and Voorhoeve (2008).

one kind of benefit is 'sufficiently' low, and the quality of another kind of benefit is 'sufficiently' high, then an outcome in which a relatively small number of people received the higher quality benefit would be better than one in which virtually any number of (otherwise) similarly situated people received the lower quality benefit.

Temkin asserts that these principles are widely endorsed. He also claims that, taken together, they show how the goodness of distributions depends on the feasible set and how transitivity may fail (pp. 42ff). He illustrates these claims with a series of spectrum cases, which conform to the following schema:

Many-Person Spectrum Case. Suppose alternative a_1 involves curing one young person of a terminal illness. Following the First Standard View, it will be possible to construct an alternative a_2 which involves curing a larger number of people of an illness which imposes slightly less harm (say, quadriplegia), so that, in a choice between a_1 and a_2 only, we intuitively ought to choose a_2 . It will also be possible to construct an alternative a_3 which cures a larger number of people than a_2 of an illness that imposes slightly less harm than the illness in a_2 , so that, in a choice between a_2 and a_3 only, we intuitively ought to choose a_3 . And so on: it will be possible to continue this sequence until we reach an alternative a_z which involves curing a very large number of people of a very minor harm, such as a headache. Any alternative a_i in this sequence ought to be chosen over the directly preceding alternative a_{i-1} when we must choose from a feasible set containing only these two alternatives. However, following the Second Standard View, from a feasible set containing only a_z and a_1 we should choose a_1 . This is so because although the number saved from the very small harm in a_z will be very large, it will not be large enough to enable these small harms to jointly outweigh the saving of one young person's life (pp. 45–52).

According to Temkin, these judgements of relative *choiceworthiness* track the all things considered relative *goodness* of the alternatives. Our intuitive judgements over this sequence are therefore that a_z is all things considered better than a_{z-1} , ..., a_i is all things considered better than a_{i-1} , ..., a_2 is all things considered better than a_1 , but a_1 is all things considered better than a_z .

What are we to make of Temkin's claims? First, the Standard Views are too vague to derive his conclusions in the Many-Person Spectrum Case. For example, without spelling out what is meant by 'sufficiently low', 'sufficiently high', 'relatively small number' and 'virtually any number', one cannot deduce from the Second Standard View the desired conclusion that a_z is worse than a_1 . (What would have been sufficient is the simple claim that curing one young person from a terminal illness is better than saving any number of people from a headache.) It would therefore have

been better if Temkin had offered a different specification of his principles, for example, by taking his lead from the careful discussion of relevant aggregative and non-aggregative principles in Fleurbaey *et al.* (2009).

Moreover, while Temkin is right that many *philosophers* share his judgements in this Spectrum Case,² he offers no evidence for his claim that these judgements are part of common-sense morality. Indeed, the only relevant study I know of finds that a *minority* of respondents endorsed both the limited form of aggregation and the limited form of non-aggregation that together imply these judgements (Cowell *et al.* 2010).

Furthermore, psychological research reveals our tendency to undervalue the importance of saving large numbers from harm (Slovic *et al.* forthcoming). One might therefore doubt whether one can trust people's intuitions about cases in which we must either save one from grave harm or save an extremely large number from very small harms (Broome 2004: 56–58).

Nonetheless, many acute thinkers agree with Temkin's judgements in the Many-Person Spectrum Case, and it is worth considering whether their judgements reflect a plausible moral view. The following is one attempt at formulating such a view.

Aggregate Relevant Claims. In the cases Temkin considers, each individual whose well-being is at stake has a claim to be helped. These people's claims compete just in case they cannot be jointly satisfied. An individual's claim is stronger the more her utility would be increased by being aided and the lower the level of utility from which this increase would take place. In addition to varying in strength, claims can be relevant or irrelevant in the face of competing claims. A claim is relevant if and only if it is sufficiently strong relative to the strongest competing claim. We should choose the alternative that satisfies the greatest sum of strength-weighted, relevant claims.

Of course, we must ask what might justify this view. I attempt such a justification elsewhere (Voorhoeve 2013; see also Kamm 1993: Ch. 8–10). Roughly, the idea is this. There is both an aggregative and non-aggregative way responding to the equal objective moral value of each individual's well-being. On the aggregative approach, satisfying N claims of a given strength is N times as important as satisfying one claim of a given strength. This approach is attractive because it asserts the unvarying and equal marginal importance of every additional person's claim of a given strength.

On the non-aggregative approach, we imaginatively take up the perspective of each person, one at a time. When we do so, we take

² See, among others, Kamm (1993: Ch. 8–10; 2007: 297–298, 484–486); Scanlon (1998: 238–241); and Otsuka (2006).

in what is at stake for each individual. After placing ourselves in each person's position, we 'step back' and take an objective perspective on the importance of what is at stake for each person, considered alone. From this non-aggregative perspective, it is most important to satisfy the strongest individual claim. Moreover, if we were instead to satisfy a different claim, then the larger the 'gap' in strength between the strongest claim and this satisfied claim, the more we would depart from what is most important. This form of concern for the claim of each person, taken alone, is a natural expression of respect for the separateness of persons.

Aggregate Relevant Claims embodies one way of partially accommodating and arbitrating between these conflicting approaches. It tells us to do what is best on the aggregative approach under the constraint that this does not depart too far from what is most important on the non-aggregative approach.

To illustrate, let us focus on three possible distributions of benefits:

death: curing one young person of a terminal illness;

considerable: saving 10,000 from a considerable impairment;

very slight: saving several billion from a very slight impairment.

Suppose that in a choice just between saving one from death and ten thousand from the considerable impairment, we ought to aggregate benefits and save the ten thousand. Suppose also that in a choice just between curing the ten thousand of the considerable impairment and several billion of the very slight impairment, we ought to save the billions. Finally, suppose that whenever it is possible to cure someone of a terminal illness, we ought *not* to aggregate the benefits of being cured of the very slight impairment.

It would follow that whether we ought to cure the billions of the very slight rather than the ten thousand of the considerable impairment will depend on whether it is feasible to cure the terminal illness. This may seem curious, but if we accept the premise that we ought not to depart 'too far' from what is most important on the non-aggregative approach, then this is perfectly reasonable. We would violate this 'relevance of claims' constraint if we were to save billions from a very slight impairment when we can instead save one from death. By assumption, we would not violate this constraint if we were to save the billions when our only other alternative is to save the ten thousand from the considerable impairment.

If this view were correct, it would vindicate Temkin's claim that the relative *choiceworthiness* of distributions depends on which other distributions are feasible. However, it is questionable whether it would also vindicate Temkin's claim that which impairment it is *better* to cure (in consequentialist terms) depends on the feasible set. Kamm (2007: 484–486), for example, suggests a nonconsequentialist interpretation

of Aggregate Relevant Claims when she writes that it would be 'disrespectful' of the one facing death for us to choose to save the billions from the very slight impairment, because it would not honour the fact that so much more is at stake for him than for anyone else. Now, on this nonconsequentialist interpretation, the transitivity of 'better than' is not threatened by the aforementioned pattern of choices in our simple example. Nor is it threatened by the choices Temkin outlines in his Many-Person Spectrum Case. For the choices in question would not always follow consequentialist betterness – one would choose to save the one from death rather than the multitude from a very minor harm on nonconsequentialist grounds.

Of course, one might instead attempt a consequentialist interpretation of Aggregate Relevant Claims. In other words, one might 'consequentialize' the view (see Broome 2004: 34–35; Brown 2011). One could hold, for example, that the violation of the non-aggregative constraint is a bad in consequentialist terms. Or one could posit that the value of satisfying the largest possible sum of strength-weighted, relevant claims always exceeds the value of satisfying irrelevant claims, no matter how many of the latter there are.

On a consequentialist interpretation, it might appear that we *do* have a violation of the transitivity of better than. However, this appearance would be misleading. Requirements of rational choice specify formal relations between values and choices. They don't tell us which substantive values to hold. Instead, they tell us that choices ought to be rational in the light of these values. In order to judge whether a moral view violates these conditions, we must therefore individuate alternatives by characteristics which that view takes to be grounds for choice (Broome 1991: Ch. 5). When we are assessing whether Aggregate Relevant Claims respects these requirements, we ought therefore to include in our description of alternatives whether, on this view, the claims in question are relevant or not. We can do this, for example, by using the superscript 'r' when the claims are relevant, and $\neg r$ when they are not. In our simple example, we then get:

From the feasible set $\{death^r; considerable^r\}$, choose *considerable^r*;

From $\{considerable^r; very\ slight^r\}$, choose *very slight^r*;

From $\{death^r; very\ slight^{\neg r}\}$, choose *death^r*;

From $\{death^r; considerable^r; very\ slight^{\neg r}\}$, choose *considerable^r*.

On the consequentialist interpretation of Aggregate Relevant Claims, these choices track the all things considered betterness of outcomes. We then have *very slight^r* is better than *considerable^r*, which is better than *death^r*, which is better than *very slight^{\neg r}*. This does not violate transitivity.

Temkin acknowledges the possibility of this response to purported violations of requirements of rational choice. He asserts, however, that to take what he calls this 'fine-grained individuation' approach is to accept that transitivity 'fails to apply':

In some cases, it *appears* that a certain relation, R , both applies to some alternatives and is intransitive. But, (...) we may decide not that the relation R is actually intransitive, but that the relation R doesn't actually apply, simpliciter, to the alternatives in question. (...) [I]t would (...) be *nontransitive* in that it would fail to apply across certain sets of alternatives to which we might have thought that it should apply (pp. 17 and 216–217, emphases in original).

It is not true that transitivity fails to apply on properly individuated alternatives. In our simple case, for example, since *considerable^r* is better than *death^r*, which is better than *very slight^r*, by transitivity, *considerable^r* is better than *very slight^r*. And since *very slight^r* is better than *considerable^r*, by transitivity, it is better than *very slight^r*. Of course, there is something uncommon about the latter judgement, since *very slight^r* cannot appear in the same feasible set as *very slight^r*. But there is nothing incoherent about the assertion that *very slight^r* realises goodness to a greater degree than *very slight^r*. Indeed, if we are to understand the view in question in a consequentialist manner, it makes perfect sense to assert that it is better to save billions from a very slight impairment when their claims are relevant than when they are not relevant.

Temkin has one further comment about this method of individuation. He writes:

perhaps the key feature that makes transitivity so fundamentally important for practical reasoning – namely, that it provides a (...) method of (...) deciding on the best of our options on the basis of a series of pairwise comparisons, where each alternative is fixed independently of the other alternatives (...) – is lost once one adopts the fine-grained solution (p. 465).

Temkin is right that Aggregate Relevant Claims may require more cumbersome decision-making than familiar principles such as utilitarianism. On the former view, whether one ought to save billions from the very slight impairment rather than ten thousand from the considerable impairment might depend on whether the feasible set includes preventing one death. By contrast, for a utilitarian, the question whether one ought to save billions from the very slight impairment rather than ten thousand from the considerable impairment can be answered without knowing the rest of the feasible set. That is convenient. But Temkin is wrong to suggest that such convenience is the key reason that transitivity is important. As mentioned above, a view which holds that a_3 is better than a_2 , a_2 is

better than a_1 , but a_1 is better than a_3 and that tells you not to choose an alternative that is strictly worse than some other alternative would rule out every choice among these three options. The singular force of the requirement that if a_3 is better than a_2 and a_2 is better than a_1 , then a_1 is *not* better than a_3 derives from the need to avoid such 'dead ends' in decision-making.

Temkin also exaggerates the inconvenience of views on which the morally relevant characteristics of the alternatives depend on our feasible set. On properly individuated alternatives, transitivity retains its pragmatic usefulness. Suppose, for example, that in Temkin's Many-Person Spectrum Case, we can save any, but only one, of all the groups from their illness. Suppose, furthermore, that when saving someone from death is an option, the claims of individuals up to and including the ninth alternative in the spectrum are relevant, but all the claims in all subsequent alternatives are irrelevant, because they are too weak relative to the competing claim to be saved from death. The first nine alternatives should then carry the superscript ' r ' and subsequent alternatives the superscript ' $\neg r$ '. It then follows that a_9^r is better than a_8^r , a_8^r is better than a_7^r , ..., a_2^r is better than a_1^r . It also follows that a_1^r is better than any of the alternatives that violate the relevance constraint, from $a_{10}^{\neg r}$ to $a_z^{\neg r}$. By transitivity, a_9^r is better than every other alternative. A series of pairwise comparisons therefore suffices to select the best alternative.

In sum, I have argued that Temkin's judgements in the Many-Person Spectrum Case do not threaten the transitivity of better than. It is a further, interesting question whether the principle that seems to underlie these judgements is correct. In Chapter 3 (and on pp. 440–5), Temkin discusses the merits of a principle (similar to Aggregate Relevant Claims) that blocks the interpersonal aggregation of very small benefits when we can instead benefit one badly off individual a great deal. His interesting discussion highlights cases in which this principle differs from familiar consequentialist views. To give an example, suppose that there are one hundred people whose utility is 1 on a scale on which 0 is equivalent to death at an early age and 100 is a long life in good health.³ We can either provide one of these people a benefit of 99 utils or all of them a benefit of 1 util. Temkin's principle favours giving it all to the person to whom we can make a large difference, but all standard consequentialist views, from leximin to utilitarianism, would favour the latter. Personally, I would here side with these familiar consequentialist views – but Temkin eloquently puts the case for the opposing view.

³ I am assuming that these numbers are elicited by the 'Standard Gamble' measure of health-related quality of life commonly used in health economics, so that they represent von Neumann–Morgenstern utilities.

2. ONE-PERSON SPECTRUM CASES

Chapters 4 and 5 cover intrapersonal tradeoffs. Temkin's core example is the following case (similar to one first formulated by Rachels [1998]):

One-Person Spectrum Case. Imagine the prospect of living for a further fixed number of years in good health, except for the fact that you will have to endure a painful episode e_1 . Imagine that this episode is two years of excruciating torture. Call this e_1 . Now imagine enduring slightly less intense torture for much longer. Call this combination of intensity of pain and the time it must be endured e_2 . It will intuitively seem better to endure e_1 . Now consider e_3 , which involves suffering a pain slightly less intense than the pain in e_2 for much longer. Again, e_2 seems better than e_3 . By iterating this kind of reasoning we can construct a sequence of alternatives e_1, e_2, \dots, e_z in which each member in this sequence is better than its successor and the final member e_z involves a pain of one extra mosquito bite per month for a very long period of time (which is, of course less or equal to the time you have left to live). But the discomfort of an extra mosquito bite per month for this very long period of time is intuitively less bad than two years of torture. So e_z is intuitively better than e_1 , violating transitivity (pp. 134–139).

To have even a moderate degree of confidence in our judgements in an example like this, it must be confined to episodes not much longer than a normal life. One reason is that for beings whose psychological states change as ours do with time, it is questionable whether the entirety of an existence that is a great deal longer than our ordinary lifespan can be the object of self-interested concern. Lewis (1983), for example, argues that a minimum of 'psychological connectedness' – being linked through memory and similarity in beliefs, values, tastes, and aims – is a necessary condition for personal identity. He also argues that a person whose psychology changes as ours does but who lives as long as Methuselah would not display a sufficient degree of psychological connectedness to remain the same person over this period. In sum, if psychological connectedness determines the boundaries of our self-concern, then for beings like us, such self-concern cannot extend far beyond our ordinary lifespan. (Of course, one can imagine beings whose psychology is different from ours – but then our confidence in our intuitive judgements about how to evaluate extended episodes of pain for them would be lowered by doubts about our ability to properly assess lives so different from ours.) This is important, because if it takes a great many steps to arrive from 'two years of torture' to 'one extra mosquito bite per month', then the sequence cannot be completed within the requisite timeframe.

In response to such concerns, Temkin suggests that it is possible to create a sequence within these confines:

A man who is 5' 2" tall is *very* short. A man who is 6' 4" is *very* tall. (...) Yet, one can move from the very short to the very tall [person] in just seven relatively small [2"] steps. If, by analogy, one could move from the pain of torture to the discomfort of a mosquito bite in seven relatively small steps, the relevant trade-off [could] be (...) between 1 year of torture and 128 years of the discomfort of a mosquito bite (p. 155).

However, Temkin does nothing to support this analogy. I therefore do not think there is much to be learned from this remark, except that it may be possible to get from being tortured by a short man to being tortured by a tall man in seven small steps.

Still, let us suppose that Temkin were to offer an adequately described sequence over which a decent sample of individuals would display intransitive preferences. Would this provide support for his critique of transitivity? I believe not. As I have argued elsewhere, psychological data suggests that we cannot trust our intuitions over such sequences (see Voorhoeve 2008, which contains relevant references). Studies show that in pairwise comparisons between two-dimensional alternatives, the dimension towards which attention is directed is often overweighted relative to the other dimension. They also show that when alternatives are similar along only one dimension, people's attention is focused on the *dissimilar* dimension and they therefore have a tendency to underweight the relative importance of the similar dimension. Between adjacent alternatives in a sequence like Temkin's, people would therefore have a tendency to underweight the importance of the intensity of pain (the similar dimension) relative to the importance of the length of time it is endured (the dissimilar dimension). Other studies on the evaluation of painful episodes show that in cases that do not direct attention to the duration of a painful episode, duration receives a surprisingly low weight; instead, pain intensity is given very great (arguably disproportionate) weight relative to duration.

What follows from these findings for the reliability of our intuitive judgements in Temkin's One-Person Spectrum? In comparing e_1 and e_z , both duration and pain intensity will capture our attention, but studies suggest our attention will primarily be directed to the great differences in pain intensity. Duration can therefore be expected to receive less weight in the choice between e_1 and e_z than it does in the choices between adjacent alternatives, where it is the focus of attention because it is the one dissimilar dimension. This explains why it may be possible to generate a sequence of the required kind. But, given the fact that both the form of similarity-based decision-making outlined here and our assessment of painful alternatives have been shown to regularly yield unreliable judgements, it also explains why our judgements about the alternatives in such a sequence ought not to be trusted.

In reply, Temkin writes:

If Voorhoeve (...) really believed that my arguments were based on similarity-based decision-making, then [he] must reject [at least one of the judgments about adjacent alternatives]. But (...) it is *deeply* implausible to do so. Given the real choice between pain of -597 for two years and pain of -596 for four years, would [he] *really* take the latter? (p. 308).

Temkin here seems to be using a cardinal scale of pain intensity on which -597 is a very intense pain and a very mild pain falls just below 0. It therefore seems that this reply ignores the constraint that the sequence must range over alternatives that are of an appreciable length. (If one were to continue the sequence as suggested by the passage – lowering the pain by 1 unit and doubling the length of time at each step – then the final alternative would involve enduring a mild pain for a time in excess of the age of the universe.) On a sequence that *does* respect this constraint and that ends up with a low-pain-intensity final alternative which seems better than the first, the differences in pain intensity between steps must be greater, so that it becomes conceivable that for at least one pair of adjacent alternatives in the sequence, in displaying a preference for the earlier alternative, one is underweighting the importance of the difference in intensity between them. In any case, I do not claim that one of the similarity-influenced judgements over pairs of adjacent alternatives in a properly specified sequence *must* be wrong – merely that studies give us reason to think that intuitive judgements over at least one of these adjacent pairs or the pair of the first and final alternatives are suspect.

In sum, Temkin asserts that ‘people’s [transitivity-violating] intuitive judgements about [the One-Person Spectrum Case] are extremely strong and not easily dislodged’ (p. 303). But he does not offer a single clear case of the kind required and provides no proper evidence of people’s responses to it. His assertion is therefore unfounded. But even if he did provide a case in which people have intransitive preferences, psychology gives us reason to doubt the reliability of their intuitions.

3. TEMKIN’S CRITICISM OF EXPECTED UTILITY THEORY

Chapter 6 contains an interesting discussion of cases in which options are not fully comparable and argues that we ought not to expect that transitivity holds when comparability does not.⁴ It also argues why liability to being ‘money pumped’ is not a sufficient reason to consider someone irrational, so that the familiar ‘money-pump argument’ in favour of transitivity fails. I am sympathetic to both points. Chapter 7 focuses on how deontic relations such as ‘permissible to do rather than’ may fail to be

⁴ Bradley (2013) offers further support for this claim and a discussion of the necessary conditions of preference-based choice without completeness.

transitive. It also contains the remarks on the 'fine-grained individuation' approach to supposed failures of transitivity to which I responded above.

Temkin realizes that one of the reasons people favour transitivity is its role in orthodox expected utility theory. In Chapter 8, he therefore challenges two of this theory's other axioms: completeness and continuity. As just noted, I am sympathetic to the idea that completeness is not rationally required, so I shall remark only on Temkin's critique of continuity. It draws on the following example from Vallentyne (1993): outcome o_1 is a long, healthy life with \$1 000 001 per year; o_2 is such a life with \$1 000 000 per year; and o_3 is a short, poverty-stricken life. If o_1 is preferred to o_2 and o_2 is preferred to o_3 , then by continuity, there is some probability $1 > p > 0$, such that a person is indifferent between o_2 for sure and a lottery $po_1 + (1-p)o_3$. But Temkin objects that it cannot be rationally required that one be prepared to trade the certainty of a good life for a lottery with a high probability of a small improvement in one's circumstances and some non-zero probability of an awful outcome.

I see the intuitive pull of Temkin's position. However, a defender of continuity has two replies available. First, the mere presence of risk may negatively affect the value of the outcomes in question. If, for example, the presence of risk in the lottery causes anxiety, this would need to be 'written into' the outcomes in question, so that the lottery does not really yield outcome o_1 with probability p , but rather o_1^- , which would be a long, healthy life with \$1 000 001 per year and a period of anxiety. If o_1^- is less good than o_1 , then it may also be less good than o_2 . It would then not merely be permissible, but rationally required to favour o_2 for sure over a lottery between the marginally inferior o_1^- and the far worse o_3 .

Second, the probability of the awful life in the lottery may be arbitrarily close to zero. There is abundant psychological evidence that our intuitive grasp of probabilities in general is tenuous. I surmise that our intuitions over alternatives involving extremely small probabilities are very unreliable. I therefore do not regard this critique of continuity as persuasive.

4. THE 'INTERNAL ASPECTS VIEW' AND THE 'ESSENTIALLY COMPARATIVE VIEW'

In Chapter 9, Temkin replies to objections to his Spectrum Arguments. I have discussed elements of these replies above. Chapters 10 to 14 contain Temkin's theoretical response to the challenges he has raised to transitivity. A key issue is the difference between the following views of the goodness of outcomes.

The Internal Aspects View. This holds that the goodness of an outcome depends only on its 'internal features', which are features that do

not depend on what other outcomes could have been brought about instead.

The Essentially Comparative View. This holds that an outcome's internal features alone do not determine its goodness, since the latter is also determined by which outcomes could have been brought about instead.

Temkin argues that transitivity holds when the Internal Aspects View is correct, but need not hold when the Essentially Comparative View applies. As my remarks so far indicate, I am sympathetic to this distinction, but not to the assertion that transitivity is threatened when the Essentially Comparative View applies. The key to retaining this distinction *and* transitivity is, I believe, to properly individuate alternatives and their associated outcomes. By way of illustration, consider again the outcomes of two properly individuated alternatives, *very slight*¹⁷ and *very slight*¹⁸, respectively. Their shared 'internal feature' is the distribution of well-being associated with the alleviation of the very slight harm. But they also each have a distinct 'essentially comparative feature' – whether the claims in question are relevant. Since they each have such a feature, we can say that their goodness is essentially comparative (assuming, that is, the aforementioned consequentialist interpretation of Aggregate Relevant Claims). This is an important conclusion, since it implies that we need to know the feasible set in order to know how to evaluate the goodness of a distribution of well-being. But transitivity holds.

5. A METHODOLOGICAL REMARK

As in his critique of transitivity and continuity, throughout the book, Temkin appeals to our intuitions about cases in which our intuitions ought not to be trusted. He asks us to contemplate cases involving trillions of people (p. 344), the possibility of sub-atomic-sized sentient beings (p. 345), and worlds with infinite numbers of people who live for eternity in Heaven or Hell (pp. 434–440). Moreover, he freely makes use of numbers to describe purportedly cardinal measures of pleasure, pain and value, without giving us an adequate sense of how these numbers are measured and therefore what they mean. Such cases are of little value for the project of evaluating distributive principles.

6. CONCLUSION

Rethinking is a creative, sometimes frustrating, but always challenging book. It contains a wealth of engaging ideas. I have focused on two: that the goodness of a distribution may be essentially comparative and that transitivity is threatened by particular case judgements. My verdict is split. I have argued that Temkin may well be right that whether

one distribution is better than another can depend on which feasible distributions it is compared to. But I have also argued that his two core cases do not threaten the transitivity of 'better than'. In the Many-Person Spectrum Case, I have argued that the apparent intransitivity disappears when we properly individuate alternatives. In the One-Person Spectrum Case, I have argued that purportedly transitivity-threatening intuitions are unreliable.

Let me conclude. In a famous soliloquy, Macbeth tries to rationalize his impending murder of King Duncan. Finding no fault with the King, he concludes: 'I have no spur // To prick the sides of my intent, but only // Vaulting ambition, which o'erleaps itself.'⁵ Temkin prosecutes his case against transitivity with greater inventiveness and verve than Macbeth can muster for his case against the King. Nonetheless, I have argued that *Rethinking* offers no good reasons for doing away with transitivity, but only vaulting intuition.

REFERENCES

- Binmore, K. and A. Voorhoeve. 2003. Defending transitivity against Zeno's paradox. *Philosophy and Public Affairs* 31: 272–279.
- Bradley, R. 2013. A note on transitivity and incompleteness. Choice Group Working Papers 9 (1). http://www2.lse.ac.uk/CPNSS/projects/CoreResearchProjects/ChoiceGroup/PDF_files/Transitive-Choice12.pdf
- Broome, J. 1991. *Weighing Goods*. Oxford: Blackwell.
- Broome, J. 2004. *Weighing Lives*. Oxford: Oxford University Press.
- Brown, C. 2011. Consequentialize this. *Ethics* 121: 749–771.
- Carlson, E. 2005. Intransitivity without Zeno's paradox. In *Recent Work on Intrinsic Value*, ed. M. J. Zimmerman and T. Rønnow-Rasmussen, 273–277. New York: Springer.
- Cowell, F., M. Fleurbaey and B. Tungodden. 2010. The tyranny puzzle in welfare economics: an empirical investigation. *STICERD working paper PEP-5*.
- Fleurbaey, M., B. Tungodden and P. Vallentyne. 2009. On the possibility of nonaggregative priority for the worst off. *Social Philosophy and Policy* 26: 258–285.
- Kamm, F. M. 1993. *Morality, Mortality, Vol. I*. Oxford: Oxford University Press.
- Kamm, F. M. 2007. *Intricate Ethics*. Oxford: Oxford University Press.
- Lewis, D. 1983. Survival and identity. In *Philosophical Papers, Vol. I*, 55–77. Oxford: Oxford University Press.
- Norcross, A. 1997. Comparing harms: headaches and human lives. *Philosophy and Public Affairs* 26: 135–167.
- Norcross, A. 1999. Intransitivity and the person-affecting principle. *Philosophy and Phenomenological Research* LIX: 769–776.
- Otsuka, M. 2006. Saving lives, moral theory, and the claims of individuals. *Philosophy and Public Affairs* 34: 110–135.
- Qizilbash, M. 2005. Transitivity and vagueness. *Economics and Philosophy* 21: 109–131.
- Rachels, S. 1998. Counterexamples to the transitivity of better than. *Australasian Journal of Philosophy* 76: 71–83.
- Scanlon, T. M. 1998. *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.

⁵ Shakespeare, *Macbeth*, Act 1, scene 7, lines 25–27.

- Slovic, P., D. Zions, A. K. Woods, R. Goodman, and D. Jinkse, forthcoming. Psychic numbing and mass atrocity. In *The Behavioral Foundations of Policy*, ed. E. Shafir. Princeton, NJ: Princeton University Press.
- Temkin, L. 1987. Intransitivity and the Mere Addition Paradox. *Philosophy and Public Affairs* 16: 138–187.
- Temkin, L. 1996. A continuum argument for intransitivity. *Philosophy and Public Affairs* 25: 175–210.
- Vallentyne, P. 1993. The connection between prudential and moral goodness. *Journal of Social Philosophy* 24: 105–128.
- Voorhoeve, A. 2008. Heuristics and biases in a purported counterexample to the acyclicity of ‘better than’. *Politics, Philosophy and Economics* 7: 285–299.
- Voorhoeve, A. 2013. How should we aggregate competing claims? Unpublished ms., available on <http://personal.lse.ac.uk/voorhoev>