

# TRANSITIVITY AND VAGUENESS

MOZAFFAR QIZILBASH

*University of East Anglia\**

---

Axiomatic utility theory plays a foundational role in some accounts of normative principles. In this context, it is sometimes argued that transitivity of “better than” is a logical truth. Larry Temkin and Stuart Rachels use various examples to argue that “better than” is non-transitive, and that transitivity is not a logical truth. These examples typically involve some sort of “discontinuity.” In his discussion of one of these examples, John Broome suggests that we should reject the claim which involves “discontinuity.” We can, I suggest, make sense of the examples which Temkin uses while sacrificing neither transitivity nor “discontinuity.” This response to Temkin’s examples involves developing and modifying James Griffin’s account of “discontinuity.” If the account of “discontinuity” seems implausible, that is because of a failure to allow for vagueness. A similar argument can be made in the context of the well-known “repugnant conclusion.”

## 1. INTRODUCTION

Axiomatic utility theory plays a foundational role in some accounts of normative principles, such as utilitarianism and contractarianism (Harsanyi 1955; Broome 1991, 1999; Binmore 1994, 1998). In this context, transitivity is sometimes assumed to hold, and is taken to be either a condition of consistency (Binmore 1994: 105) or a truth of logic (Broome 1991: 11–12). The standard consistency axioms have recently been much criticized (Sugden 1985; Anand 1993; Sen 1995 *inter alia*). John Broome

\* This paper emerged from a discussion in one of John Broome’s seminars. I am very grateful to John Broome, Erik Carlson, James Griffin, Wlodek Rabinowicz, Stuart Rachels, and Larry Temkin for allowing me to read various forthcoming manuscripts which are mentioned in the paper. I am also extremely grateful to Gustaf Arrhenius, Walter Bossert, Luc Bovens, John Broome, Richard Cookson, Robin Cubitt, James Griffin, Graham Loomes, Ben McQuillin, Shepley Orr, Wlodek Rabinowicz, Stuart Rachels, Bob Sugden, and two anonymous referees for comments on earlier versions. Any errors or omissions are mine.

has defended the axioms of expected utility theory in the light of some of these criticisms (Broome 1991, 2004). Broome claims that “better than” is a transitive relation, and that this is a matter of logic.

Larry Temkin (1987, 1994, 1996, 2000) and Stuart Rachels (1993, 1998, 2001, forthcoming) have argued, to the contrary, that “better than” (or “all things considered better than”) is not transitive. In this paper, I shall focus on Temkin’s version of this argument, though the conclusions I draw from the discussion of Temkin’s work are also relevant to Rachels’ version of the argument, which is based on very similar examples. While Rachels and Temkin claim that “better than” is not transitive, they also make the weaker claim that transitivity of “better than” is not a truth of logic. Undermining the *status* of transitivity of “better than” is central to their purpose. Rachels and Temkin make these claims using specific examples. One of the distinctive characteristics of these examples is that they involve a series of alternatives where adjacent members of the series are slightly different from each other in some respect. These examples are closely related to cases involving small differences which W. E. Armstrong (1939) and Amos Tversky (1969) used in discussions of non-transitivity. Armstrong’s examples are often related to a Sorites paradox (Fishburn 1988) because they involve small differences. This paradox is closely associated with, and much discussed in the growing literature on, vague predicates (Keefe and Smith 1996). Yet in spite of a growing literature on imprecise or “fuzzy” preferences in economics (Barrett and Pattanaik 1989; Dubourg *et al.*, 1994; Dasgupta and Deb 1996 *inter alia*) discussion of vague predicates and their implications for economics has been limited.

In this paper, I focus on Temkin’s version of the argument and argue in favour of one line of response to Temkin’s examples, which invokes James Griffin’s notion of “discontinuity.” I suggest that Temkin’s worry about this response is best dealt with by modifying this line of response by allowing for vagueness. In making the argument, I use a “supervaluationist” account of vagueness. Before making my central argument, I explain John Broome’s recent discussion, which involves a very clear formalisation of one of Temkin’s examples. These examples are closely related to examples which are discussed in the context of population issues. I argue that, in that context, allowing for vagueness can illuminate discussions of the well-known “repugnant conclusion.”

The paper is structured as follows: in section 2, I explain Temkin’s examples; in section 3, I discuss Broome’s response to one of these examples; in section 4, I discuss Griffin’s notion of “discontinuity” and how it relates to Temkin’s examples and to Broome’s discussion; in section 5, I show how the “discontinuity” view can be modified using an account of vagueness; in section 6, I relate the discussion to the “repugnant conclusion”; and section 7 concludes.

## 2. TEMKIN'S EXAMPLES

While Larry Temkin and Stuart Rachels have both argued that “better than” is not transitive, the precise examples they use are slightly different. Temkin’s examples are, nonetheless, variations of examples which Stuart Rachels developed earlier (Rachels 1993), so that Rachels must take much of the credit for them. Since Temkin changed Rachels’ formulation of the example which follows, I focus on Temkin’s discussion.<sup>1</sup> This example – henceforth, the “hangnail example” – invokes three claims: (1) “for any unpleasant or ‘negative’ experience, no matter what the intensity and duration of that experience, it would be better to have that experience than one that was only a little less intense but twice as long”; (2) “there is a continuum of unpleasant or ‘negative’ experiences ranging in intensity, for example, from extreme forms of torture to the mild discomfort of a hangnail”; and (3) “a mild discomfort for the duration of one’s life would be preferable to two years of excruciating torture, no matter the length of one’s life” (Temkin 1996: 179). Here Temkin uses “preferable to” to mean “better than.” Furthermore, he does not mean “continuum” to imply that the range of negative experiences is infinitely variable. Rather his claim is that there is a discrete range of experiences with some a bit better or worse than others. Since Temkin’s use of “continuum” may have misled some commentators (such as Binmore and Voorhoeve 2003a), I adopt the term “spectrum” which he also uses to capture the idea of a discrete range of levels.<sup>2</sup> Temkin thinks that if we accept these three claims – and he suspects that most of us will accept them – we must doubt, or reject, the transitivity of “better than.”

The transitivity of some relation  $R$  requires that for any  $x$ ,  $y$ , and  $z$ , if  $xRy$  and  $yRz$  then  $xRz$ . Why do claims (1)–(3) lead us to reject the transitivity of “better than”? One version of Temkin’s argument involves various lengthy lives  $A$ ,  $B$ ,  $C$ ,  $D$ , etc. Suppose that  $A$  contains two years of excruciating torture, while  $B$  contains four years of torture which is slightly less excruciating than the torture in  $A$ . Then imagine another lengthy life  $C$  which has eight years of torture which is slightly less painful than the torture in  $B$ . If we accept claim (2) there is a spectrum of painful experiences that allows us to form this sort of sequence of lengthy lives. The sequence would go from lives  $A$ ,  $B$ ,  $C$  to other lengthy lives  $D$  – with 16 years of pain slightly less extreme than in  $C$  –  $E$ ,  $F$  and so on, till we arrive at  $Z$ , a lengthy life with only mild discomfort for a very long period. If claim (1) is true then:  $A$  is better than  $B$ ,  $B$  better than  $C$ , and so on, so that  $Y$  is better

<sup>1</sup> I do not discuss Temkin’s (1987) related work on the “mere addition paradox”, though section 5 is closely related to Temkin’s subject in that paper.

<sup>2</sup> Arrhenius (2000: 162–64) uses the term “discreteness” to capture this sort of “spectrum” notion in population problems. It is easy to see how confusion has abounded in related discussions of “continuity” in this literature.

than Z. Transitivity then implies that A is better than Z. However, if claim (3) is true then Z is better than A. So we have a contradiction.<sup>3</sup>

Temkin takes this to be a counterexample to transitivity of “better than.” He suggests that failures of transitivity can arise because different factors are relevant in making different sorts of comparisons. Similar sorts of factors may, for example, be relevant in making comparisons between lives which are close to each other in the sequence – such as pairwise comparisons between A and B, B and C etc. – while others might matter in making comparisons between lives which are further from each other in the sequence – such as A and Z. Furthermore, Temkin suggests that some factors may be more significant in some comparisons than in others (Temkin 1996: 193). He thinks that, even if this example is not a “knockdown” counterexample to the transitivity of “better than” – since he accepts that there are ways of responding to it which preserve transitivity – there is no way of responding to it which does not involve sacrificing some other plausible belief (Temkin 1996: 210). Even if transitivity of “better than” is still accepted, furthermore, transitivity may seem far from self-evident.<sup>4</sup> There are, of course, different ways of responding to this example. On the one hand, one can object to any of the three claims that lead to non-transitivity. On the other hand, one might claim that the three claims do not lead to non-transitivity. Ken Binmore and Alex Voorhoeve (2003a) have adopted the second approach using a specific utility function,<sup>5</sup> though their response assumes that Temkin is adopting some notion of “continuity” rather than merely a spectrum of discrete levels (see Carlson 2003).

Claim (3) involves what James Griffin has termed “discontinuity.” Griffin’s discussion in *Well-Being* focuses on “prudential values” – the elements of a good life. In the cases of “discontinuity” that Griffin (1986: 85) discusses the focus is on comparisons between certain pairs of values or goods. In these comparisons, there comes a point where however much more of one value or good one might have, that cannot outweigh a certain amount of the other value or good.<sup>6</sup> This sort of case

<sup>3</sup> There are various ways of objecting to the argument, but my purpose is only to explain it.

<sup>4</sup> The argument for non-transitivity in this example is related to classic discussions in the literatures in economics and psychology (Armstrong 1939; Tversky 1969) which relate to non-transitivity of “preferred or indifferent to”. It is also related to apparent non-transitivities or violations of consistency axioms in cases where there are multiple factors involved in choices (Sugden 1985; Anand 1993; Sen 1995).

<sup>5</sup> Binmore and Voorhoeve (2003a: 274) use a utility function with two arguments,  $p$  and  $t$ , where  $p$  is the level of discomfort and  $t$  is the length of time that it is endured, and the utility function is  $u(p, t) = -pt/(1 + t)$ . Binmore and Voorhoeve (2003a: 277–78) also show that a similar utility function can be used to deal with a related example close to Griffin’s “beauty example” which I discuss in sections 3 and 4.

<sup>6</sup> Temkin has himself noted that his arguments are related to Griffin’s (Temkin 1996: 190, and forthcoming), and he has made the link with Griffin’s discussion more explicitly in a

is sometimes treated as a form of “incommensurability” which relates to the suspension of “trade-offs” (where a “trade-off” involves loss in one value or good being compensated by a gain in another) between different values or goods (Anderson 1993).<sup>7</sup> Temkin’s third claim involves this sort of “discontinuity” in terms of the value “pain avoidance.” In Temkin’s example, the idea is that any period of mild discomfort is better than some period of torture, once the torture is painful enough.

Recently, Temkin has made much the same argument using an example involving illnesses. I shall call it the “illness example” for short. One version of a “standard view,” which he thinks most of us would readily accept, involves a form of “discontinuity” – when this is understood in terms of a suspension of “trade-offs.” He calls this the “second standard view”: “[t]rade-offs between quality and quantity are sometimes undesirable: if the difference in the severity of illnesses is sufficiently large, it would be better to prevent or cure a small number of people from suffering the severer illness, than any number of people from the less severe illness” (Temkin forthcoming: 4). The examples of more severe illnesses which Temkin mentions include AIDS, quadriplegia, and severe psychosis. Examples of less severe illnesses he mentions include a minor nosebleed, a slight cold, and a short mild headache.

Temkin also thinks that, other things being equal, numbers matter, especially when comparing illnesses which are sufficiently close in severity. So he thinks that another claim – which he calls the “first standard view” – would be readily accepted by most people. It is the view that: “it is better to prevent or cure a larger number of people from suffering a less severe illness, than a smaller number of people from a more severe illness, if the number of people aided is sufficiently greater, and if the differences in severity are not too great” (Temkin forthcoming: 4). As with claims (1) and (3), these two “standard views” express “ground-level” judgements, in the sense that they do not follow from any other judgements that Temkin explicitly mentions. The thought is, presumably, that these views capture important intuitions with which any ethical theory should be compatible. However, if we accept these standard views and there is

paper about abortion (Temkin 2000: 264–65). Arrhenius (2000: 92–93) associates this notion of “discontinuity” with the idea of “lexical priority” while Rachels (2001: 219–25) uses the notions of “lexicity” and “lexical priority” and Crisp (1997: 30–31) invokes the same notion in explaining John Stuart Mill’s distinction between “higher” and “lower” pleasures. Crisp (1992) and Ryberg (1996) discuss “discontinuity” in the context of the repugnant conclusion.

<sup>7</sup> Economists usually analyse such a suspension of trade-offs in terms of “lexicographical preferences”. Such preferences violate the standard axiom of “continuity.” The continuity axiom is most commonly defined in consumer theory in terms of consumption bundles: for any consumption bundle, the set of bundles which is at least as good as that bundle, and the set of bundles which that bundle is at least as good as, are closed.

a “spectrum” of actual or potential illnesses moving from very severe illnesses progressively to less severe illnesses in small discrete steps, we are again, according to Temkin, led to non-transitivity (Temkin forthcoming: 13–14). Temkin argues that we must either reject one of the two standard views, or, alternatively, we must give up the transitivity of “better than.” The first standard view plays the role of claim (1), and the second standard view plays the role of claim (3) in the hangnail example.

Claim (2), which is invoked in the hangnail example, is related, in this context, to the requirement that there is a discrete range of illnesses such that adjacent illnesses in the range are “sufficiently close” in their severity (Temkin forthcoming: 13). The requirement is not precise, because the notion of “sufficiently close” is not precise. The assumption underlying Temkin’s argument seems to be that a discrete change in the severity of an illness is “sufficiently small” if and only if the change in severity can always be weighed against a change in the number of people saved from each illness when comparing options which involve saving people from these illnesses. If this is so, saving some number of people from one illness is worse than saving some larger number of people from some less severe illness, when the difference in severity of illnesses is “sufficiently small.” However, this is a very specialised sense of “sufficiently small.” I term Temkin’s assumption the “spectrum requirement”. In the illness example it involves two joint claims: (i) there is a discrete range of illnesses, and adjacent illnesses in the range are “sufficiently close” and; (ii) illnesses are “sufficiently close” as regards severity if the difference in severity does not undermine the possibility of weighing changes in severity against changes in the number of people cured in evaluating states of affairs. This requirement is clearly not supposed to apply exclusively to the range of illnesses we actually encounter – it applies also to any illnesses we can imagine. However, if one defines “sufficiently close” in this way and the spectrum requirement is met, the first standard view holds for the range of actual or potential illnesses one can think of.

### 3. JOHN BROOME’S RESPONSE TO THE ILLNESS EXAMPLE

John Broome has responded to the illness example. In his “A comment on Temkin’s trade-offs” (Broome forthcoming), Broome formalises the two “standard views” in Temkin’s discussion in such a way that they jointly lead to non-transitivity of “better than.” In making sense of Temkin’s health care example, Broome treats a “small number” as any number smaller than 10. “I” stands for “illness” and illnesses are indexed,  $I_1, I_2, I_3, \dots$  where  $I_1$  is more severe than  $I_2$ ,  $I_2$  is more severe than  $I_3$  and so on. No assumption is made about how much more severe  $I_2$  ( $I_3$  etc.) is than  $I_1$  ( $I_2$  etc.): all that is assumed is that there is a discrete range of illnesses which can be ranked

in terms of “more severe than.”<sup>8</sup> The number of people saved from  $I_1$  ( $I_2$  etc.) are similarly indexed and written:  $n_1$  ( $n_2$ , etc.). Broome then states Temkin’s second standard view as follows:

SV2: There are two illnesses  $I_1$  and  $I_k$  and a number  $n_1$  less than 10 such that, for any number  $n_k$ , curing  $n_1$  people of  $I_1$  is better than curing  $n_k$  people of  $I_k$ .

Broome’s version of Temkin’s first standard view runs as follows:

SV1\*: For every pair of illnesses  $I_1$  and  $I_k$  and every number  $n_1$  less than 10, there is a sequence of illnesses  $I_1, \dots, I_k$  and a sequence of numbers  $n_1, \dots, n_k$  such that for every  $I_j$  and  $I_{(j+1)}$  in the sequence, curing  $n_{(j+1)}$  people of  $I_{(j+1)}$  is better than curing  $n_j$  people of  $I_j$ .

In combination, the validity of these views leads to non-transitivity of “better than” on Broome’s account. Broome thinks that there are theoretical grounds – which he does not elaborate on in his paper – for accepting SV1\*. So, given that he thinks that “better than” is transitive, Broome argues that Temkin ought to reject SV2. It is worth noting that the spectrum requirement is not invoked in either SV1\* or SV2, though in Temkin’s arguments that requirement is central. The “discontinuity” claim in Temkin’s illness example is embodied in SV2. So Broome’s suggestion seems to be that we should reject that part of Temkin’s argument which involves “discontinuity.”

In *Weighing Lives* Broome (2004) presents a slightly different version of this argument. In this version, Broome sets  $n_1$  equal to 1. He, in effect, treats SV1\* and SV2 as properties which a sequence of illnesses might have. A sequence of illnesses  $I_1, I_2, \dots, I_k$  may or may not have the following two properties:<sup>9</sup>

*Property 1:* curing a single person of  $I_1$  is better than curing any number of people, however large, of  $I_k$ .

and,

<sup>8</sup> While there are no explicit definitions of “illness” or “severity” in this literature, one can take as given here some definition of “illness” in terms of a departure from some norm of health or functioning. “Severity” of an illness can involve a number of dimensions – such as the amount of pain involved, the threat to life posed by the illness etc. Given that it may be hard to rank some illnesses in terms of severity, it need not be assumed that all illnesses can be ranked in terms of severity: all that is needed is a discrete range of actual or potential illnesses which can be ranked.

<sup>9</sup> Broome (2004: 56) uses slightly different terminology referring to property 1 as “property (a)” and property 2 as “property (b).”

*Property 2:* there is a sequence of numbers  $1, n_2, \dots, n_k$ , starting with 1, with the property that curing  $n_i$  people of  $I_i$  is better than curing  $n_{(i-1)}$  people of  $I_{(i-1)}$ , for any value of  $i$  from 2 to  $k$ .

In making sense of Broome's response to Temkin, I shall write "better than" as " $>$ ". I treat it as a primitive relation.<sup>10</sup> I also use the following symbols: "&" for conjunction; " $\neg$ " for negation; " $\in$ " for "is a member of"; " $\rightarrow$ " for the conditional "if ... then"; " $\leftrightarrow$ " for "if and only if"; and " $\forall$ " for the universal quantifier ("for all"). I write the set of "options" – where "options" are actual or potential objects of choice – as  $\mathbf{X}$ .  $>$  may have the following properties: *transitivity*:  $\forall x, y, z \in \mathbf{X}, (x > y \ \& \ y > z) \rightarrow x > z$ ; *asymmetry*:  $\forall x, y \in \mathbf{X}, x > y \rightarrow \neg(y > x)$ ; and *irreflexivity*:  $\forall x \in \mathbf{X}, \neg(x > x)$ . While  $>$  has the second and third of these properties as a matter of logic, I shall suspend any judgement about the logical truth of transitivity. I treat transitivity of  $>$  as a working proposition used in proving theorems. In particular, I assume that the following postulate holds:

*The Betterness Postulate:*  $>$  has the properties of irreflexivity, asymmetry, and transitivity.

For the moment, the only sort of options considered involve curing people from actual or conceivable illnesses. I shall write the option "curing  $n_i$  people of  $I_i$ " as " $C(n_i, I_i)$ ." One can express Broome's version of Temkin's view as follows: if some finite sequence of illnesses  $I_1, I_2, \dots, I_k$  has properties 1 and 2, then  $>$  is not transitive. This proposition is best understood in the light of the following simple theorem:

**THEOREM 1.** If a finite sequence of illnesses  $I_1, I_2, \dots, I_k$  has property 2, it does not have property 1.

**PROOF.** Suppose that this is not so. Suppose that the sequence of illnesses  $I_1, I_2, \dots, I_k$  has property 2, as well as property 1. If  $I_1, I_2, \dots, I_k$  has property 2, then it follows that there is a sequence of numbers,  $1, n_2, \dots, n_k$  such that  $C(n_2, I_2) > C(1, I_1)$ ;  $C(n_3, I_3) > C(n_2, I_2), \dots$ , &  $C(n_k, I_k) > C(n_{(k-1)}, I_{(k-1)})$ . It then follows from the betterness postulate (transitivity of  $>$ ) that  $C(n_k, I_k) > C(1, I_1)$ . But the sequence  $I_1, I_2, \dots, I_k$  has property 1, and  $C(1, I_1) > C(n_k, I_k)$ . By the betterness postulate (asymmetry of  $>$ ), this implies  $\neg[C(n_k, I_k) > C(1, I_1)]$ . We have a contradiction. ■

I prove this result because, while Temkin's argument for non-transitivity remains controversial, this theorem is not. The theorem makes the tension between properties 1 and 2, which is at the heart of Broome's discussion,

<sup>10</sup> The analysis would not be much altered if  $>$  were defined in terms of "at least as good as", and some minimal assumptions were made.



explicit. What can one conclude from this theorem? We might conclude that: the sequence does not have property 1; the sequence does not have property 2; or we might conclude that both properties 1 and 2 are plausible and that transitivity should be rejected. This third conclusion is the one which Temkin effectively favours, and this leads to what I have called “Broome’s version of Temkin’s view.” For those who accept that transitivity of  $>$  is a logical truth, the first two conclusions are more appealing. Broome (2004: 56) favours rejecting the idea that any sequence of the sort  $I_1, I_2, \dots, I_k$  has property 1. He thinks that there are sequences that have property 2 and he thinks that the error in Temkin’s argument is the view that there is some illness such that curing one person, or a small number of people, of that illness is better than curing *any* number of people of a more minor illness. That is, he thinks, just an intuition involving very large numbers. Broome (2004: 56–59) argues that our intuitions about very large numbers are shaky and that we ought not to drop transitivity on the basis of such intuitions. He thus questions the intuition behind property 1. However, property 1 does seem plausible to many people. Broome only makes the case explicitly in the context of Temkin’s discussion of illnesses, though it could equally be made in the hangnail example as well as Rachels’ examples. In the hangnail example, again, it would be the “discontinuity” claim – claim (3) – which Broome would reject, and some would not follow him.

Theorem 1 can be established without invoking transitivity of  $>$ , if  $\mathbf{X}$  is a finite set. It holds if  $>$  is “acyclic”:

Acyclicity of  $>$ :  $>$  is acyclic over a finite  $\mathbf{X} \leftrightarrow \forall a, b, \dots, z \in \mathbf{X}: [(a > b, b > c, \dots, y > z) \rightarrow \neg (z > a)]$ .

Given that Temkin’s argument focuses on a discrete range of illnesses, it is harmless to assume that the set of options is finite. If  $>$  is acyclic, a shorter proof of theorem 1 is available. If  $I_1, I_2, \dots, I_k$  has property 2 then  $C(n_2, I_2) > C(1, I_1)$ ;  $C(n_3, I_3) > C(n_2, I_2)$ ;  $\dots$ ; &  $C(n_k, I_k) > C(n_{k-1}, I_{k-1})$  (as in the above proof). Clearly, acyclicity of  $>$  implies  $\neg[C(1, I_1) > C(n_k, I_k)]$ . If  $I_1, I_2, \dots, I_k$  has property 1 then  $C(1, I_1) > C(n_k, I_k)$ . So we have a contradiction.

This result is not surprising. Many examples in the economics literature which are proposed as counterexamples to transitivity and other consistency axioms are also thought to be counterexamples to acyclicity (Sugden 1985; Anand 1993). Furthermore, the key issue that Temkin has raised is about the *status* of the transitivity axiom, not about the minimal conditions under which theorem 1 can be proven. If we really think that the intuitions behind properties 1 and 2 are strong enough, we might decide to reject both transitivity *and* acyclicity of  $>$  if they imply that one of the two properties must be rejected. Certainly, there is a presumption in Temkin’s

work that the intuitions in his examples are strong enough to be on a par with any argument in favour of transitivity of  $>$ . This presumption is contentious. Nonetheless, theorem 1 suggests that properties 1 and 2 are in conflict, and the fact that we can prove it without invoking the betterness postulate means that the conflict does not depend on assuming transitivity. I shall argue that in this conflict there is a case for sacrificing property 2 rather than property 1. This involves accepting the “discontinuity” claim which Broome rejects.

#### 4. GRIFFIN’S “DISCONTINUITY” VIEW AND TEMKIN’S EXAMPLES

If one accepts that there are “discontinuities” in these examples, without rejecting transitivity of  $>$ , then, on Temkin’s account, one must accept that one of the other claims in Temkin’s arguments is false. As regards “Broome’s version of Temkin’s view,” one must reject property 2, if one accepts property 1 as well as transitivity. In terms of Temkin’s discussion of the illness example, if one accepts Temkin’s “discontinuity” claim (standard view 2) as well as transitivity, one must give up either: (a) the first standard view; or (b) the spectrum requirement. To do so one needs to find a view that can make sense of this example without accepting (a) or (b). In terms of Broome’s discussion, this view must clarify why one should reject property 2.

We thus need to find a plausible view which involves small changes of the sort used in these examples, and is compatible with the idea that a small difference in severity can undermine trade-offs. One candidate view is embodied in James Griffin’s discussion of “discontinuity.” In a well-known passage in *Well-Being* Griffin focuses on an example involving a specific form of enjoyment – the appreciation of beauty. I shall call it the “beauty example.” Griffin (1986: 86) thinks that 50 years of life involving the appreciation of the very best Dutch paintings – the very best Rembrandts, Vermeers, and de Hoochs – might be outranked by 55 years at a slightly lower level – involving the rest of the Dutch school, but not the very best of it; and 55 years at this lower level might be outranked by 60 years at a level of artistic appreciation which is slightly lower still. While he thinks that this “step-by-step approach” seems irresistible, he resists it because he thinks that as one repeatedly reduces the quantity of enjoyment in the level of appreciation, and enough beauty is lost, there may come a point where there is no beauty left to appreciate, and beauty might be replaced with the “kicks of kitsch . . . [which] are different” (Griffin 1986: 87). The “discontinuity” claim in this example is: no amount of years lived at a quality of life which merely involves the enjoyment constituted by the kicks of kitsch outranks some number of years lived at a quality of life which involves the appreciation of beauty.

This example involves a substantive difference between different kinds of enjoyment. It is not meant to be a formal account of “discontinuity.” Nonetheless, one might derive a formal account based on it as follows. We need to distinguish “beautiful” objects of appreciation from those objects of appreciation of which is enjoyable, but which are not beautiful. To do this, consider an alternative which involves the most enjoyable form of appreciation one can think of for some finite duration of time. If we reduce the quality of enjoyment involved in the appreciation a little, and increase the duration of the enjoyment enough, we may, nonetheless, be faced with a new alternative which is better. If we continue this process, on the “discontinuity” account, after some point it will no longer be the case that there is an alternative which is better than the previous alternative. On this account, then, all the objects in this sequence up to this point are objects of “beauty,” and all objects which follow it are not beautiful (they may be “kitsch” or fail to be “beautiful” in some other way). On Griffin’s account the substantive distinction between beauty and kitsch must fit with this formal account for the beauty example to work.

It is easy to see how this logic can be applied to the hangnail and illness examples. For illustrative purposes, I shall focus on the illness example. If one were to think about the illness example in these terms, then one might judge that some illnesses are “serious.” One might think of these – in substantive terms – as illnesses which can undermine the ability to lead a flourishing life. These would include the more severe illnesses that Temkin mentions, such as AIDS, quadriplegia and so on. Illnesses which are not of this sort – “non-serious illnesses” – include the less severe illnesses which Temkin mentions: a minor nosebleed and a mild cold. An alternative, formal, way of defining the difference between the serious and non-serious would go like this. First consider the most severe actual illness one can think of and consider the option of saving some number of people of such an illness. This is a serious illness. It is plausible that saving a larger number of people from a slightly less severe illness is better than saving a smaller number of people from the more severe illness. If so, the less severe illness is still serious. Next, one might consider the option of saving the larger number from the less severe illness and go through the same exercise – finding an illness which is a little less severe still, and a new option involving saving some even larger number of people from the less severe of the two illnesses. If the new option is better, then the less severe of the two illnesses is still serious. Yet, it may be that, as we reduce the severity of illnesses in this step-by-step fashion, there is a point at which it is not true that we can find a new option such that it is better than the previous one in the sequence. All illnesses which are less severe than this point are “non-serious,” and illnesses at or above this level of severity are “serious.” The substantive and formal distinctions between “serious” and “non-serious” illnesses may or may not coincide.

On the formal version of the “discontinuity” view just developed, one of Temkin’s claims ought to be rejected. Indeed, when we reach the “last” serious illness, it is no longer true that a small reduction in the severity of the illness is “sufficiently small” in the sense which is relevant to the “spectrum requirement.” That is, it is no longer true that one can trade-off a small change in severity against the number of people saved from the various illnesses. So the “spectrum requirement” is not met. It is also easy to see (or to check) that, in terms of Broome’s version of Temkin’s view, this “discontinuity” account involves rejecting property 2 for some set (or sets) of illnesses. So on this account one could, without contradiction, accept property 1 and transitivity.

The way in which the “discontinuity” view has been developed here involves different kinds of illnesses (objects of appreciation etc.). Temkin has responded to this variety of objection to his position, in the context of the hangnail example. He writes that

[m]any people worry that my counterexample trades on the fact that the pain of intense torture seems to be of an entirely different kind than the pain of the hangnail. This is why no amount of the latter can outweigh two years of the former. This observation is, I think, right on target. But it does not undermine my argument. (Temkin 1996: 194)

The reason Temkin does not think that the observation undermines his argument is this. Consider the spectrum of negative experiences ranging from extreme forms of torture to the pain of the hangnail. On his view “[i]n comparison with a small amount of torture, a hangnail’s duration basically does not matter” (Temkin 1996: 194). Temkin thinks that the duration of the negative experience plays a more significant role in comparisons between experiences which differ only in degree rather than in kind. It is easy to see why this argument runs contrary to Griffin’s position as it has just been spelt out. The idea is that a difference in “degree” – that is, a small difference in the quantity of pain in this example – cannot amount to a difference in kind. Yet it is just this which, I suggest, is implied in Griffin’s rejection of the “step-by-step approach” in the beauty example. In that example, a small reduction in the quantity of enjoyment can, as we saw, make the difference between an object of beauty and mere kitsch. The implication of the “discontinuity” view, as it is developed here, in the context of the hangnail example, would be that a difference in “degree” *can* amount to a difference in kind.<sup>11</sup>

<sup>11</sup> Indeed, Griffin (forthcoming: 3) has made this point in discussing “discontinuities” in a recent response to Temkin, stating that “[a]t some stage, differences in quantity produce differences in quality.” So the way in which Griffin’s views are developed here is consistent with Griffin’s published views in this respect. A similar result – to the effect that a small change in quantity can lead to a change in quality – can be found in the literature on “Millian

In fact, Temkin goes on to discuss this line of objection directly in the context of the hangnail example. He thinks that: “[f]rom the fact that a large number of differences in degree can amount to a difference in kind, it surely does not follow that there must be some point where a small difference in degree itself results in a difference in kind” (Temkin 1996: 196). He goes further and argues that this line of objection “would insist that if we believe that torture’s pain is of a different kind than a hangnail’s, then there *must* be some pain which lies between that of torture and that of a hangnail, such that it would be *worse* to have a pain of that intensity for some duration than a pain which was *a little less intense* for twice (or three, or five times as long). I find this deeply implausible” (Temkin 1996: 197). Yet what Temkin finds implausible is the precise implication of Griffin’s account of “discontinuity” as it is interpreted here. The line of objection which Temkin finds implausible might, nonetheless, seem less implausible if one allows for vagueness as well as “discontinuity.”

### 5. VAGUENESS AND “DISCONTINUITY”

In the beauty example, Griffin rejects what he calls the “step-by-step” approach. In this example the steps involved are: 50 years of life involving appreciation of the very best of the Dutch school; 55 years at a quality of life at a slightly lower level; and so on. He rejects the “step-by-step” approach because it:

presents us with two embarrassments: a Sorites paradox and a slippery slope. If we take enough pebbles away from a heap, it ceases to be a heap. But since one pebble more or less could never make the difference between its being a heap or not, if we remove the pebbles singly it can never cease to be a heap. Similarly one might argue, with “appreciation of beauty” or “deep loving relations”; with slow step-by-step changes we can never lose the appreciation of beauty or deep loving relations. But we obviously can. (Griffin 1986: 87)

If there is a danger of a Sorites paradox involved in this example, clearly the same worry must arise for the examples Temkin presents, which have the same structure. Consider the illness example. In that example, if one accepts the “discontinuity” account I just developed, some illnesses are “serious.” It seems implausible that a slight change in the severity of an illness can make for a difference in the kind of illness it is. As we saw, Temkin makes this sort of claim – that a difference in “degree” cannot make for a difference in “kind” – in his discussion of the hangnail example. If a similar claim is true in the context of the illness example, then repeatedly reducing the severity of a serious illness a little, step-by-step, ought to

discontinuities” (Arrhenius and Rabinowicz 2003) and in the discussion of Ryberg’s views on higher and lower pleasures (Rabinowicz forthcoming).

leave us with an illness of the same kind: a serious illness. No small change in itself can make for a difference in kind. However, that must be false. If severity is sufficiently reduced in this step-by-step fashion one will be left with a non-serious illness. This is a standard example of a Sorites paradox.

If these examples can lead to Sorites paradoxes, it is quite plausible that they involve vague predicates. Typical examples of such predicates include "tall" and "bald." In their excellent introduction to the philosophical literature on vague predicates, Rosanna Keefe and Peter Smith tell us that susceptibility to Sorites paradoxes is one of the classic characteristics of vague predicates (Keefe and Smith 1996: 3). The two other characteristics of vague predicates which Keefe and Smith single out are: that they involve borderline cases – cases where it is not clear whether the predicate applies; and that they involve rough borderlines. So in the case of the predicate "tall," there may be people who count as "borderline tall" so that they are neither clearly tall, nor clearly not tall. Furthermore, there is no sharp borderline between those who are tall and those who are not tall. The predicate "tall" only involves one dimension (height). However, multidimensionality is also sometimes thought to be relevant to vagueness. For example, in the case of another vague predicate – "nice" – there may be more than one dimension which is relevant to judging whether someone is nice or not. The fact the someone does well on some dimensions and less well on others, may contribute to her being borderline nice (Keefe and Smith 1996: 5).

Temkin (1996: 197–202) and Rachels (forthcoming: 18–20) have both carefully argued that their claims cannot be dismissed in the way that Sorites paradoxes are rejected.<sup>12</sup> Suppose, then, that one *rejects* the claim that there is any Sorites paradox in these examples. If one follows the "discontinuity" view, one might still hold that the borderline between

<sup>12</sup> Temkin (1996: 199–201) suggests – in the context of an argument leading to a Sorites paradox – that the conclusion that someone is both hairy and not hairy involves a key premise which might be interpreted so that a *small* difference is treated as *no* difference. He argues that in his argument he does not make this sort of error. There are clear differences between the outcomes being compared in the examples he discusses. (Rachels forthcoming:19, makes a similar point). In an alternative interpretation of the relevant premise, Temkin argues that the relations "no significant difference between" or "similar to" play a role – and that these are non-transitive relations. By contrast, Temkin states that his argument does not involve these relations, but only the relation "better than." Temkin argues that problems involving Sorites paradoxes can be dealt with by a "tightening of our proposed linguistic usage" but that nothing comparable is available for his proposed counterexample(s) to transitivity. Note, however, that in the example just discussed the issue I have focused on is about whether a small difference in *quantity* ("degree") – of severity of illnesses, or of pain, or of appreciation – can amount to a change in the nature of illnesses, or pain, or of appreciation ("kind"). In this context, it certainly is true – on Temkin's account – that a small difference in quantity, a difference of "degree," makes *no* difference as regards "kind."

objects of different kinds is imprecise. So one might accept that the line between serious and non-serious illnesses may be rough, in the illness example. Equally, some illnesses may be considered borderline serious. Similarly, in the beauty example the borderline between objects of beauty and kitsch may not be precise and some objects may be borderline beautiful. So these examples may involve vagueness, even if we put the issue of whether they involve Sorites paradoxes to one side. Finally, multi-dimensionality is clearly relevant to Temkin's examples. As we saw earlier, he himself tries to "explain" the non-transitivity that arises in these cases in terms of different factors being relevant to, or having different levels of significance in, different pairwise comparisons. It may then be illuminating to analyse these examples using some account of vague predicates.

There are different views of vagueness. These views are usually classified into three rough categories: epistemic views; degree theories; and supervaluationist views. On epistemic views, there is actually a sharp borderline, even in the case of vague predicates. It is just that we do not and *cannot* know where this borderline is. The chief weakness of this view is that it is often not clear that our failure to pin down a sharp borderline in the case of vague predicates has to do with any constraint on knowledge. Degree theories, by contrast, allow for a rough borderline. According to such theories, in the case of "tall," there would be cases where it is definitely true that someone is tall, cases where it is definitely true that a person is not tall and yet other cases where it is true *to some degree* that a person is tall. The chief weakness of this view involves the notion of a degree of truth (Keefe and Smith 1996: 46–49). One way of developing the notion of a degree of truth goes like this. If Jim is taller than John, it is *more true* that Jim is tall than it is true that John is tall. Yet if both Jim and John are very tall, this is very implausible. Rather it is definitely true that both are tall. This objection renders some degree theories unattractive. Finally, there are supervaluationist views. These suppose that the truth of statements involving vague predicates depends on how they are made more precise, or – in alternative terminology – on how they are "precisified" or "sharpened". In the best-known version of supervaluationism, Kit Fine's version, a vague sentence is "super-true" if it is true on all "admissible" ways of making it more precise or "sharpening" it (Fine 1996: 132; Keefe and Smith 1996: 23). A certain amount of residual ("higher-order") vagueness remains, because "admissible" is vague. One advantage of the supervaluationist view is that it does not, like the epistemic view, suppose that there is any single sharp boundary. Furthermore, each "sharpening" of a borderline is arbitrary.<sup>13</sup> Another

<sup>13</sup> In the remainder of the paper, when I use the notion of making a borderline "more precise," or "precise" or "sharpening" a borderline, I mean making it as precise as possible. In Fine's terms, this involves a "complete" specification of a "base point" (Fine 1996: 126).

advantage is that it does not require any notion of a degree of truth. There are, of course, also criticisms of supervenience (Keefe and Smith 1996: 32–35). Some worry, for example, that it is inappropriate for an account of vagueness to be based on the idea of making vague statements more precise. Nonetheless, I shall use Fine's supervenience account to modify Griffin's "discontinuity" view here.

To see how a supervenience analysis of Temkin's example would go, I focus on the illness example. If "serious" is vague, there is no sharp transition between illnesses that are serious and those that are not. A supervenience account of this example might then go as follows. There are different "admissible" ways of sharpening the borderline between serious and non-serious illnesses. There may be some illness,  $I_j$ , which is "serious" on some admissible way (or ways) of making the predicate "serious" more precise, but "non-serious" on another way (or other ways) of making the borderline precise. On Fine's account, " $I_j$  is serious" is "super-true" if and only if " $I_j$  is serious" is true for all admissible ways of making "serious" more precise. In cases where " $I_j$  is serious" is true on some admissible way(s) of making "serious" more precise but not on another (or others), " $I_j$  is serious" is not super-true.  $I_j$  then falls in a vague zone of illnesses between those which are *definitely* serious – i.e., illnesses for which it is super-true that they are serious – and those which are *definitely* non-serious – i.e. illnesses for which it is super-true that they are not serious.<sup>14</sup> For each sharpening of "serious" in the vague zone, the "discontinuity" view would nonetheless be true. That is, for each sharpening, a small change in severity would make for a change in kind. To this degree, Griffin's "discontinuity" view is consistent with the supervenience account of this example.

Nonetheless, allowing for vagueness using this supervenience account does involve modifying the "discontinuity" view in a significant manner. This is because the supervenience account leaves us with a vague zone between those illnesses which are definitely serious and those which are definitely not serious. Consider again a sequence of illnesses ordered according to their severity. Suppose also that there is a precise set of admissible ways of sharpening the borderline between serious and non-serious illnesses and that  $I_s$  is *definitely* serious and that all illnesses between (and including)  $I_{(s+1)}$  and  $I_{(s+10)}$  are neither definitely serious nor definitely non-serious.  $I_{(s+11)}$  is, nonetheless, *definitely* non-serious. A small reduction in the severity of  $I_s$  will not leave us with an illness which is definitely non-serious. Instead, it will leave us with a new illness which is neither definitely serious nor definitely non-serious. There is thus no abrupt transition from illnesses which are definitely serious to those which are

<sup>14</sup> Fine's supervenience also allows for "higher-order vagueness" by allowing for vagueness of "admissible."



definitely non-serious.<sup>15</sup> There is, furthermore, no *single point* where there is a change from being definitely serious to being definitely non-serious. Put another way, there is no unique point where a difference in “degree” can amount to an *unambiguous* difference in “kind.” This is one way of addressing the key worry that Temkin had about the line of objection to the hangnail example which was discussed at the end of section 4. When one allows for vagueness in this way, the modified “discontinuity” view is, I suggest, plausible.<sup>16</sup>

What light does this analysis shed on the formal distinction between “serious” and “non-serious” illnesses given in section 4? Recall that the distinction was formally defined by starting with the most severe illness one could think of and considering an option involving saving some number of people from suffering from that illness. We then compare this option with another involving saving some larger number of people from an illness which is a little less severe. If we find some larger number such that saving that number of people from the less severe illness is better than curing the smaller number of people from the more severe illness, then the less severe illness is “serious.” Now consider the example that I have just presented. We have supposed that  $I_s$  is serious on this way of defining “serious.” So beginning with the most severe illness one can think of, we have, in steps, reached an alternative involving saving some number of people,  $n_s$ , of  $I_s$  which is better than saving some smaller number of people from an illness which is a little more severe. The discussion implies that it is not definitely true that there is some number  $n_{(s+1)}$  such that it is better to save that number of people of  $I_{(s+1)}$  than to save  $n_s$  people from  $I_s$ . A similar claim can be made for any other illness in the vague zone. Nonetheless, there is definitely no  $n_{(s+11)}$  such that it is better to save that number of people of  $I_{(s+11)}$  than it is to save  $n_s$  people from  $I_s$ .

Vagueness clearly does imply that in some cases we cannot make definite evaluative judgements. An inability to make such judgements can clearly lead to indeterminacy in practical reasoning, if one takes a “teleological” view so that what we ought to do depends only on the

<sup>15</sup> In the example I am discussing, I am making the assumption that there is an exact set of admissible sharpenings. The argument can, however, be made even if we relax this assumption, and allow for “higher-order vagueness.”

<sup>16</sup> Ken Binmore and Alex Voorhoeve (2003b) analyse Warren Quinn’s torture example (Quinn 1990). Quinn’s example is importantly different to Temkin’s because it involves a sequence with *imperceptible* differences in one dimension – involving comfort or discomfort – alongside perceptible differences in another dimension. Binmore and Voorhoeve allow for “indeterminacy” about the transition from comfort to discomfort – but they analyse this in terms of *uncertainty* about whether one is in comfort or discomfort. They can, thus, analyse the example using expected utility theory. The key contrast between the vagueness account adopted here and Binmore and Voorhoeve’s approach is that here there is vagueness in the *language* used to describe the world, whereas in Binmore and Voorhoeve’s approach the indeterminacy arises from uncertainty about the state of the *world*.

goodness of the alternatives. One obvious way to avoid such indeterminacy is to fix on some particular sharpening of “serious” and to work with that. Governments often do this, in dealing with vagueness. In the case of the vague predicate “poor,” for example, a single poverty line which separates the poor from the non-poor is often selected for the purposes of government policy. That is a rather arbitrary manner to deal with vagueness. It may, on occasion, be the only practical option. There are, nonetheless, other ways of making judgements in some cases. For example, one could treat an illness as “serious” for practical purposes, even when it is in the vague zone, if it is serious on enough (or a majority) or sharpenings of “serious”.

I have focused my analysis on a supervaluationist account of vagueness. It is worth noting that the epistemic view gives an alternative analysis of this example. If we adopt that view, there *is* an unknowable sharp borderline between illnesses which are serious and those which are not. On this view, if “serious” is vague, there would in reality be an abrupt transition from serious to non-serious illnesses. Vagueness about the point of transition from serious to non-serious illnesses is just ignorance. Given our ignorance about the location of this point, we need to leave a margin for error about its location (Williamson 1996: 279). This epistemic account preserves Griffin’s thought that there is a unique point where a change of degree makes for a change in kind. I find this account implausible, because of the presupposition that there are actually sharp borderlines even when predicates are vague.

## 6. THE REPUGNANT CONCLUSION

The central characteristics of the key examples discussed above are present in several well-known problems that both economists and philosophers have puzzled over. I shall focus in this section on well-known population puzzles. While Temkin has discussed some of these in related work on non-transitivity (Temkin 1987), the links between the examples just discussed and population puzzles is most clear in Stuart Rachels’ recent works (notably his 2001 and forthcoming). I focus here on Derek Parfit’s (Parfit 1984: 388) “repugnant conclusion” which has been much discussed (Dasgupta 1993; Blackorby *et al.*, 1997; Broome 2004 *inter alia*). It runs as follows:

*The Repugnant Conclusion:* for any possible population of at least 10 billion people, all with a very high quality of life, there must be some much larger imaginable population whose existence, if other things are equal, would be better, even though its members have lives that are barely worth living.

One might be led to this conclusion if one supposes that something like property 2 holds for levels of well-being. To see this point, suppose

that levels of well-being are written “ $w$ ” and are indexed so that  $w_1$  is a very high quality of life,  $w_2$  is less high than  $w_1$ ,  $w_3$  is less high than  $w_2$ , and so on.  $w_k$  is a level of well-being where life is only barely worth living. There is a finite range of levels of well-being between  $w_1$  and  $w_k$ . I shall again use  $n$  to refer to numbers, where  $n_2$  ( $n_3$ ) is smaller than  $n_3$  ( $n_4$ ) and so on. One can then define the following property that a sequence of levels of well-being might have:

*Property 3:* there is a sequence of numbers 10 billion,  $n_2, n_3, \dots, n_k$ , starting with 10 billion, with the property that a population of  $n_i$  people living at level of well-being  $w_i$  is better than a population of  $n_{(i-1)}$  living at  $w_{(i-1)}$ , for any integer  $i$  such that  $k \geq i > 10$  billion.

We can also define a property analogous to property 1. I shall call this property 4:

*Property 4:* a population of 10 billion people living at well-being level  $w_1$  is better than any number,  $n_k$ , of people living at level of well-being  $w_k$ .

This property embodies the relevant “discontinuity” claim in the population context, though the choice of the number 10 billion is made for convenience here. If we write a state of affairs with a population of  $n_i$  people living at level  $w_i$  as  $P(n_i, w_i)$ , we can prove a theorem which is very similar to theorem 1. It is:

**THEOREM 2.** If a sequence  $w_1, \dots, w_k$  has property 3, then it does not have property 4.

**PROOF.** Suppose to the contrary that the sequence of illnesses  $w_1, w_2, \dots, w_k$  has property 3 as well as property 4. If  $w_1, w_2, \dots, w_k$  has property 3 then it follows that there is a sequence of numbers, 10 billion,  $n_2, \dots, n_k$  such that  $P(n_2, w_2) > P(10 \text{ billion}, w_1)$ ;  $P(n_3, w_3) > P(n_2, w_2), \dots, P(n_k, w_k) > P(n_{(k-1)}, w_{(k-1)})$ . It then follows from the betterness postulate (transitivity) that  $P(n_k, w_k) > P(10 \text{ billion}, w_1)$ . But the sequence  $w_1, w_2, \dots, w_k$  has property 4, so that  $P(10 \text{ billion}, w_1) > P(n_k, w_k)$ . The betterness postulate (asymmetry) now implies  $\neg[P(n_k, w_k) > P(10 \text{ billion}, w_1)]$ . We have a contradiction. ■

One can respond to this result by: rejecting property 4; rejecting property 3; or rejecting transitivity of  $>$ . Rejecting property 4 may not be the right response. If we accept that property 3 holds for any sequence of levels of well-being, and we hold onto the betterness postulate, while rejecting property 4, it is clear from the proof that we must conclude that  $P(n_k, w_k) > P(10 \text{ billion}, w_1)$  and accept the repugnant conclusion. A similar analysis has led Stuart Rachels (forthcoming) to argue that one has a choice between accepting the repugnant conclusion and rejecting transitivity

of  $>$ . Both Temkin and Rachels would reject transitivity of  $>$  in this context. Yet clearly we also have the option of rejecting property 3, which is similar to property 2. One could then accept transitivity and reject the repugnant conclusion.<sup>17</sup>

If one follows the logic of section 4, then, rather than rejecting property 4, one might develop Griffin's "discontinuity" view instead. One way to do this would be to suppose that some levels of well-being constitute "satisfactory" lives. One might fill out this notion using a substantive account. So a satisfactory life might be one which contains all, or enough, of the things that make a life go better. A life below this level may be of such low quality that it could not count as a life of dignity. Alternatively, one might define a satisfactory life more formally. To check whether or not a level of well-being constitutes a satisfactory life we would go through the same sort of formal test as before. Suppose that  $w_1$  is the highest actual quality of life one can think of. A life at this level of well-being is satisfactory. Now consider a situation where 10 million people are living at this level of well-being. Next suppose that  $w_2$  is a slightly lower level of well-being. If there is some number  $n_2$  such that  $P(n_2, w_2) > P(10 \text{ billion}, w_1)$   $w_2$  is a level of well-being which constitutes a satisfactory life. We then look for a slightly lower level of well-being  $w_3$  and repeat the exercise, i.e., we check whether  $P(n_3, w_3) > P(n_2, w_2)$ . Once we find a level of well-being which does not classify as "a satisfactory life" using this test, we classify it as "not satisfactory." As long as there is some level of well-being in the sequence  $w_1, \dots, w_k$  which is not a satisfactory life then, clearly, the sequence will not have property 3. As long as a life which is barely worth living is not satisfactory, on this account, one can reject the repugnant conclusion while accepting transitivity.

Again it is plausible that there is no precise transition from levels of living which constitute a satisfactory life to those which do not. In this context, again, one can argue that the borderline between lives which are, and are not, satisfactory is imprecise. There is, rather, a range of levels of well-being which neither definitely constitute satisfactory lives, nor definitely constitute non-satisfactory lives. These fall in a vague zone between those lives which are definitely satisfactory and those which are definitely not satisfactory. This argument is closely related to an argument which John Broome has made about vagueness about the level of

<sup>17</sup> Temkin claims that those who respond in the way that Broome does to the illness example – by suggesting that our intuitions about very large numbers are shaky – will end up accepting the repugnant conclusion. Broome writes that "Temkin creates the impression that, if you deny SV2, you will find yourself committed to what Parfit calls the 'repugnant conclusion'... the impression is false" (Broome forthcoming: 3). Broome is right about this. However, Temkin clearly has in mind something which is *similar to*, but not the same as, SV2: property 4. One can easily object to the intuitions which lead one to hold property 4 on the grounds that they involve very large numbers.

well-being above (below) which adding more people to the population is a good (bad) thing (Broome 2004). So while the argument of this paper differs from Broome's position with respect to the analysis of Temkin's views, the suggestion that allowing for vagueness can help us to understand population problems leads to a position that is very like Broome's own view in the context of some population problems.

In a much discussed endnote, James Griffin (1986: 338–40) has famously argued that his arguments about "discontinuity" might be applied to population problems. In this note, Griffin argues that, beginning with a very high quality of life, as we constantly reduce the level of well-being in small steps we might reach a point where "people's capacity to appreciate beauty, to form deep loving relationships, to accomplish something with their lives beyond just staying alive... all disappear" (Griffin 1986: 340). Griffin seems to be trying, in this passage, to find the right point, using his substantive account of well-being where there is – in the terms I have been using – a transition from a life which is "satisfactory" to one which is not. Parfit (1986: 160–64) has discussed a related position (involving "perfectionism"). Griffin's and Parfit's attempts to develop such a substantive account of the lives involved in these examples have been criticised by Jesper Ryberg (1996), who raises important worries for these attempts to develop the "discontinuity" view in the population context. Griffin's discussion of this point was, of course, meant to be no more than suggestive (coming as it does in an endnote).

In the version of this view that I have just outlined, however, we do not need a substantive account of the nature of lives which are "satisfactory" and those which are barely worth living. The search for that account will, no doubt, go on. For the purposes of the argument, all we need is the formal test just given. The key point, whether one uses a substantive or a formal account, is to accept that there may be a vague borderline between those lives which are, and those which are not, satisfactory. One can then more plausibly hold onto the "discontinuity" claim (embodied in property 4), while rejecting property 3, accepting transitivity of  $>$  and rejecting the repugnant conclusion.

## 7. CONCLUSIONS

This paper has focused on some of Larry Temkin's "counterexamples" to the transitivity of  $>$ . Griffin's "discontinuity" view is interpreted and developed so that it involves a point where there is a transition from one kind of object (illness, experience or object of appreciation) to another. Since it is implausible that there is such a single point, the "discontinuity" view leads to an implausible conclusion. Temkin rejects one line of objection to his view for just this reason. However, there are good reasons for believing that Temkin's examples are best understood when analysed in terms of

vague predicates. When one allows for vagueness, it is no longer the case that there is a sharp transition from one kind of object to another. That makes a modified version of the “discontinuity” view plausible. A similar modified “discontinuity” argument is developed in the context of the “repugnant conclusion.”

## REFERENCES

- Anderson, E. 1993. *Value in ethics and economics*. Cambridge University Press
- Armstrong, W. 1939. The determinateness of the utility function. *Economic Journal* 49:453–67
- Arrhenius, G. 2000. *Future generations: a challenge for moral theory*. University of Uppsala
- Arrhenius, G. and W. Rabinowicz. 2003. On Millian discontinuities. Typescript
- Barrett, C. and P. Pattanaik. 1989. Fuzzy sets, preference and choice: some conceptual issues. *Bulletin of Economic Research* 41:229–53
- Binmore, K. 1994. *Game theory and the social contract. Volume 1: playing fair*. MIT Press
- Binmore, K. 1998. *Game theory and the social contract. Volume 2: just playing*. MIT Press
- Binmore, K. and A. Voorhoeve. 2003a. Defending transitivity against Zeno’s Paradox. *Philosophy and Public Affairs* 31:272–79
- Binmore, K. and A. Voorhoeve. 2003b. Defending transitivity from the Sorites Paradox, Typescript. University College London, Department of Economics
- Blackorby, C., W. Bossert, and D. Donaldson. 1997. Critical-level utilitarianism and the population-ethics dilemma. *Economics and Philosophy* 13:197–230
- Broome, J. 1991. *Weighing goods: equality, uncertainty and time*. Blackwell
- Broome, J. 2004. *Weighing lives*. Oxford: Oxford University Press
- Broome, J. forthcoming. A comment on Temkin’s trade-offs. In *Ethical issues in human resource allocation*, ed. D. Wikler and C. Murray. World Health Organization
- Carlson, E. 2003. On some recent examples of intransitive betterness. In *Logic, law, morality*, ed. R. Sliwinski: 181–95. Department of Philosophy, Uppsala University
- Crisp, R. 1992. Utilitarianism and the life of virtue. *The Philosophical Quarterly* 42:139–60
- Crisp, R. 1997. *Mill on Utilitarianism*. Routledge
- Dasgupta, P. 1993. *An inquiry into well-being and destitution*. Oxford University Press
- Dasgupta, M. and R. Deb. 1996. Transitivity and fuzzy preferences. *Social Choice and Welfare* 13: 305–18
- Dubourg, W., M. Jones-Lee, and G. Loomes. 1994. Imprecise preferences and the WTP-WTA disparity. *Journal of Risk and Uncertainty* 9:115–33
- Fine, K. 1996. Vagueness, truth and logic. In *Vagueness: a reader*, ed. R. Keefe and P. Smith 119–50. MIT Press
- Fishburn, P. 1988. *Nonlinear preference and utility theory*. Harvester Wheatsheaf
- Griffin, J. 1986. *Well-being: its meaning, measurement and moral importance*. Clarendon Press
- Griffin, J. forthcoming. Larry Temkin, Health distribution and the problem of trade-offs. In *Ethical issues in human resource allocation*, ed. D. Wikler and C. Murray. World Health Organization
- Harsanyi, J. 1955. Cardinal welfare, individualistic ethics and interpersonal comparisons of utility. *Journal of Political Economy* 63:309–21
- Keefe, R. and P. Smith. 1996. Introduction: theories of vagueness. In *Vagueness: a reader*, ed. R. Keefe and P. Smith: 1–56. MIT Press
- Parfit, D. 1984. *Reasons and persons*. Oxford University Press
- Parfit, D. 1986. Overpopulation and the quality of life. In *Applied ethics*, ed. P. Singer: 145–64. Oxford University Press
- Quinn, W. 1990. The puzzle of the self-torturer. *Philosophical Studies* 59:79–90
- Rabinowicz, W. forthcoming. Ryberg’s doubts about higher and lower pleasures – put to rest? *Ethical Theory and Moral Practice*

- Rachels, S. 1993. *A theory of beneficence*. Unpublished undergraduate thesis. University of Oxford
- Rachels, S. 1998. Counterexamples to the transitivity of 'better than'. *Australasian Journal of Philosophy* 12:71–83
- Rachels, S. 2001. A set of solutions to Parfit's problems. *Noûs* 32:214–38
- Rachels, S. forthcoming. Repugnance or intransitivity: a repugnant but forced choice. To appear in an anthology edited by J. Ryberg and T. Tannsjo. Kluwer
- Ryberg, J. 1996. Parfit's repugnant conclusion. *Philosophical Quarterly* 46:202–13
- Sen, A. 1995. Is purely internal consistency of choice bizarre? In *World, mind and ethics. Essays on the ethical philosophy of Bernard Williams*, ed. J. Altham and R. Harrison, 19–31. Cambridge University Press
- Sugden, R. 1985. Why be consistent? A critical analysis of consistency requirements in choice theory. *Economica* 52:167–83
- Temkin, L. 1987. Intransitivity and the mere addition paradox. *Philosophy and Public Affairs* 16:138–87
- Temkin, L. 1994. Weighing goods: some questions and comments. *Philosophy and Public Affairs* 23:350–80
- Temkin, L. 1996. A continuum argument for intransitivity. *Philosophy and Public Affairs* 25: 175–210
- Temkin, L. 2000. An abortion argument and the threat of intransitivity. In *Well-being and morality: essays in honour of James Griffin*, ed. R. Crisp and B. Hooker: 263–79. Oxford University Press
- Temkin, L. forthcoming. Aggregation and problems about trade-offs. In *Ethical issues in human resource allocation*, ed. D. Wikler and C. Murray. World Health Organization
- Tversky, A. 1969. Intransitivity of preferences. *Psychological Review* 76:31–48
- Williamson, T. 1996. Vagueness and ignorance. In *Vagueness: a reader*, ed. R. Keefe and P. Smith: 265–80