Defending the transitivity of "better than" in the face of axiological relativity - draft

Irene Bosco

August 12, 2020

At first sight, if A is all-things-considered better than B and B is all-things-considered better than C, the judgment that A is all-things-considered better than C seems to follow. This should be a straightforward consequence of the transitivity of the "all-things-considered better than" relation. However, if we deny that transitivity is a logical axiom of those relations involving comparatives, then it might turn out that betterness is not transitive. Following Temkin's terminology[1], a relation R will be defined as *nontransitive* iff either it is intransitive (i.e. aRb and bRc yet $\neg(aRc)$) or if transitivity cannot be applied, for there is no single relation R that holds between all the alternatives. More specifically, this could happen in a case where the obtaining of the better-than relation depends on different factors according to the nature of the relata.

As a result, there are two main strategies available to question the axiom of transitivity: the first is by providing an effective counterexample displaying cRa as well as $(aRb \wedge bRc)$, hence yielding a cycle. The second, more general way, is to show that all-things-considered betterness is a relation that conceptually involves comparison-dependent betterness relations, so that it is *likely* that transitivity fails to apply across different alternatives. The aim of this paper is to investigate part of the recent attack on the transitivity of better than, to conclude that there seems to be no cogent reason why this would not be a transitive relation.

The first section provides some background for the discussion, namely an analysis of the betterness relation in terms of goodness. The second section deals with what grounds goodness itself: good-making properties and factors that could change according to the objects at hand. The paradigm counterexamples to the transitivity of better than, i.e. Spectrum Arguments, will be briefly addressed in §3. The focus of §4 is Temkin's positive arguments for the nontransitivity of better than by introducing the Internal Aspects View and the Essentially Comparative one. Section §5 is devoted to the underlying reasons for the introduction of ECV, that is to say person-affecting considerations. Section §6 concludes.

---

[1]This term is first introduced in Temkin (2012, p.16-17)

1 The Betterness relation

The meaning of "better than"

Before we start asking which logical properties a relation has, it is key to have a clear understanding of the meaning of that relation, viz. at least an outline of the truth-conditions in virtue of which the relation obtains. Looking at the meaning of comparative forms for other gradable adjectives, I think a natural pre-theoretical understanding of the meaning of a claim like "A is better than B" is simply that the goodness of A somehow exceeds that of B,[2] just like "A is heavier than B" is true if and only if A's weight is greater than B's[3]. This view presupposes three related but distinct claims, which will be discussed in turn:

(i) both A and B have some sets of properties which ground their respective goodness

(ii) these goodnesses are comparable and there is a difference in their respective levels

(iii) A's goodness level ranks higher than B's

(i)

Clearly, the choice of which properties are to be taken into account depends on what sense of goodness one is tracking and on which properties the relata could be bearers of. Lots of things of different kinds could be compared and judged better than another of a sufficiently similar kind, and the properties that are relevant to make claims like "L.A. is better than N.Y.C.", "Tolstoy is better than Dostoevskij" and "lying is better than killing" are different ones. Moreover, even between entities of the same kind, we often need to specify a smaller subset of possible properties that are relevant for the comparison, for example by adding that "Tolstoy is better than Dostoevskij *at writing poems*". This means that Tolstoy's specific goodness level of the type limited to poetry exceeds Dostoevskij's. Given that the meaning of goodness is context-sensitive, this sort of additions is vital if we want "better than" to be an intelligible, unambiguous and informative relation. Formally, one can add a further entity c, so that "better than" becomes a three-place relation, relating not only two objects but also the concept which specifies the

---

[2]Defining "better" in terms of goodness is the route taken by Temkin (2012, p.369), as well as a view supported by other commentators like Huemer (2013); I find taking goodness as the primitive notion more explanatory-effective than taking as fundamental the notion of "better than". Note that these views are not be confused with that which gives an account of "good" in terms of "better than a given context-relative standard level"

[3]This simple view can be found in linguistics analyses of comparatives using different approaches, like that based on degrees (Kennedy, 2007) and on tropes (Moltmann, 2009)

type of goodness[4] involved in the comparison, thus the sets of relevant properties. Properties relevant to comparing poets are different from those relevant to compare novelists or chess players. "A is c-better than B" means that there is a set of properties of A and there is a set of properties of B such that the respective c-goodness supervenes on them. Note that so far nothing has been said regarding whether the relevant properties are intrinsic or not; also, it might be the case that the two sets comprise properties that are different in kind, so that the relevant properties of A might be of a different sort than those of B. As we shall see, this fact is problematic if a property could be realized by multiple sets of factors and it's not clear how these combine to make up the same type of property; unfortunately, more often than not this is the case. Lastly, if the objects to be compared are very distinct, then no set of properties could give rise to a goodness of the given type. For instance, the "morally-better than" relation between the number 7 and Lake Michigan clearly doesn't obtain, for neither could be the bearer of a morally relevant property, but nor can the relation "being a better vacation destination", for that sort of goodness hardly can be grounded by properties of numbers. For now, we should limit ourselves to conclude that for any concept c, whether a c-better relation between two entities A and B obtains or not depends on the c-goodness-making features instantiated by A and B respectively.

(ii)

Once assumed that both A and B are bearers of the same kind of c-goodness (e.g. goodness as ability to play tennis, as being a weekend plan, as being suited for a job), then it must be the case that the respective extents of this property as it is instantiated by each object are comparable. This follows if one endorses a "Comparability Thesis" as advanced by Dorr et al. (2020), according to which either of A or B is at least as c-good as the other. Consequently, to claim that A is (strictly) better than B, one must register a difference in A's and B's goodness levels.

Moreover, to spot the difference we need to find a way to compare one goodness against the other. Sometimes it's easier to measure the extent of the same property instantiated by two objects by using a common scale: we are able to say whether A is heavier than B by comparing their weights, i.e. the ratio between the extent of the specific property and an appropriate unit of measurement. Other times it is possible to do a direct comparison, for example by putting A and B side-by-side and observe who is taller. In this case, there's no use of a measuring function that assigns a level to each height with reference to an external object of a fixed height. The same reasoning applies to higher-order properties like goodness,

---

[4]Note that relevant types of properties are associated to concepts like "moral goodness", "height", "beauty", "health", etc., hence all those properties ascribable using a gradable adjective

which depends on factors that in turn could be directly or indirectly compared vis-à-vis those of the same kind.

(iii)

Finally, note that comparability plus a difference in the respective two degrees of goodness implies a form of orderability. Hence, the meaning of the c-better-than relation requires that the goodness of A and that of B are associated to different levels, on which we can establish a (total) order, according to a specific binary relation[5]. If A's goodness level precedes B's, then A is c-better than B. As a further consequence, we can conclude that the c-better-than relation is irreflexive and asymmetric[6]: it cannot be the case that A is c-better than itself, and if A is c-better than B, then the same relation does not obtain between B and A. In passing, it is worth pointing out that if a relation R is irreflexive yet not antisymmetric, it follows that R is not transitive. Hence, having established that both properties must hold for better than, a set of sufficient conditions for nontransitivity already fails.

Betterness in action: on cakes, Mozart and large people

To better illustrate the points made above, it could be helpful to look at some practical examples. First of all, we might wonder whether the experience of eating a piece of Sachertorte is better than the experience of listening to Mozart's Overture from "The marriage of Figaro". Here there could be several types of goodness that one could be interested in, like good for one's health or good as a discussion topic with musicians; suppose we are interested in evaluating the two experiences regardless of instrumental considerations, that is assessing their intrinsic goodness, cashed out, on a hedonist account, as pleasure for a certain unspecified individual. Arguably, each experience is different both at a neuronal and a phenomenological level, so that the respective pleasantness character is realized by different states and factors[7]. Since we lack a thermometer by which measuring and comparing the degree associated to the pleasantness of each state, nor can we compare the two goodnesses vis-á-vis, the only choice left is to turn

---

[5]For example, two different levels of height could be ordered using the common $>$ relation; other properties might not be represented on a numeric scale, so that an order on classes needs to be envisaged. For any specific c-better relation, the ordering relation can be chosen arbitrarily; however, once it has been defined it becomes part of the conceptual meaning of that relation, to the point that switching to a different ordering function might not preserve the identity of the relation itself

[6]A relation R on a set S is irreflexive if and only if for any $a, b \in S$, if aRb then $a \neq b$; R is said to be asymmetric if from aRb then $\neg(aRb)$

[7]Here I am assuming a multiple-realizability view of mental states; opponents of this view might deny that the same type of pleasantness could be a higher-order property of both experiences. In this case, I think one might compare the goodness, meant as pleasantness, of eating a piece of two different cakes or listening to two different compositions. The point of the example is to show that the set of properties that constitutes the supervenience base are different, and I think this could be the case even if the objects being compared are more similar in kind

to "pleasure-tracking" factors. For example, following Mill one could consider what an ideal subject who has experienced both pleasures would prefer. However plausible, suppose that thanks to technological advances, some neuronal activities associated with pleasure become easily detectable and measurable. Now, in order to establish which of the two experience is better we would face the difficult task of weighing the subject's preference against what her brain scan says: there might not be a straightforward way of doing it.

When more than just one dimension is relevant and it's not possible to directly compare the property at hand, we must be careful in deciding how to aggregate the different factors. This clearly applies to nonnormative concepts as well: Temkin (2012, p.164) advances an example concerning the relation "being larger than", for which the considerations for "better than" given above apply as well[8]. Since we cannot directly compare if a person is larger than another, we need to use dominance principles for the relevant largeness-making features, which, the example goes, are a person's height and weight. To begin with, there's the implicit commitment to the view that the higher-order property of "largeness" for people has a supervenience base comprised of the properties of height and weight. Moreover, according to Temkin's account, it is sufficient that A is higher or heavier than B to affirm that A is larger than B. Since the two disjuncts are not mutually exclusive, namely if A is heavier than B it could be the case that B is taller than A, the larger-than relation so defined is not asymmetric. As "larger than" is also irreflexive, for nothing could be taller or heavier than itself, it follows that it is not transitive.

Jettisoning asymmetry could be an easy way to obtain a failure of transitivity, yet this move leads to a notion that, pace Temkin, is conceptually incoherent. The only possibility to make sense of a statement like "A is larger than B and B is larger than A" is that the total orders imposed on A's and B's largeness levels in each conjunct are described by distinct ordering relations. My view is that this makes the first-conjunct larger-than relation different in meaning from the second: either we are dealing with an incoherent concept or we are equivocating on two distinct relations. To be sure, the different orderings on the largeness levels do not depend on the set of properties that grounds largeness, for the same properties, i.e. the value of one's height and weight, would ground the "smallness property" as well[9]; however, the choice if an ordering function is part of the meaning of the largeness. That said, from a logical point of

[8]In fact, my idea is that any relation based on comparative forms like "being F-er/more F than" could be translated into a relation involving the relevant F-notion of goodness, e.g. "being better at being F", where the good-tracking properties would be exactly those that ground F

[9]For example, it is possible to define the relation "shorter than" exactly by using the same considerations on heights as for "taller than" but employing a different order, viz. $<$, to obtain that the height that ranks higher is actually that of the shorter person

view it is still possible to retain asymmetry and at the same time derive intransitivity[10]; whether it is also conceptually justified needs to be seen.

2   Multidimensional concepts - handle with care

The sense of goodness we are interested in is the "moral goodness", namely the ethical value associated to a given state of affairs or outcome. The relevant better-than relation is thus meant to capture which of two states of affairs is morally better[11]. Obviously, whether a certain extent of moral goodness exceeds another is not something we can come to know by putting two states of affairs side by side. Similarly, it is not measurable by comparing it to a fixed measuring unit; the associated ethical value of a state of affairs has to be derived indirectly, by weighing and combining morally-good relevant factors. As it is the case for other senses of goodness or more in general for complex properties, moral goodness is a multidimensional concept. This means that there are multiple goodness-tracking factors, and moreover that different sets of them constitute a supervenience base for the moral-goodness property.

Take the example of beauty in a person. This concept depends on different criteria also according to a person's age or gender. There could be some common factors, say having proportionate facial features, but others, like height, are weighted differently, and some, e.g. having a thick beard, are not relevant for some types of people whatsoever. Despite these differences, I believe that comparisons like "Adam is more beautiful than Brianna" or "the father is more beautiful than his son" are still meaningful. Hence, what really matters here seems to be the ability of both relata to be bearers of a property of the same type, so that the two token-properties are comparable. Arguably, there are cases where although both objects could be ascribable the property of being beautiful, a comparison sounds odd. "Brianna is more beautiful than Lake Michigan" can hardly be taken literally, for even if both a person and a lake can beautiful, comparing their respective beauty levels seems impossible.

I am not sure what the reason for this failure is. A possible diagnosis could be that the above is a polysemy case, plausible for the grounding sets of factors of each relatum are likely disjoint. For the same reason, sentences like "this morning's math test was harder than a piece of quartz" are liable to a metaphorical reading only. The adjective "hard" has a different meaning when the bearers of the associated property are minerals instead of math tests, and the factors that ground the hardness of the former are entirely different from those that ground the hardness of the latter. However, at the end of this section we will

---

[10]An example is the "more ertnog" relation as envisaged by Handfield (2016)
[11]From now on, unless when specified, the relation "better than" is to be understood in the moral sense

discuss a counterexample to the view that tracking different factors is a necessary condition for being two conceptually different relations.

All things considered

If we claim that A is better than B all things considered, we are committing ourselves to a judgment which explicitly concerns the (moral) goodness of two states of affairs that takes into account all and only their respective good-making features. The supervenience recipe could thus change accordingly, and since this sort of recipes are often unknown, the mere task of determining the set of maximally relevant factors to instantiate moral goodness is an extremely challenging one.

When dealing with populations, morally relevant factors could be some statistics concerning the well-being of the people involved in each of the alternative states of affairs. There is a standard procedure for producing a population axiology that is - in theory - fairly straightforward. Assumed that it is possible to measure a person's well-being on a numerical scale via a utility function, then all that is required is a further function that takes as argument an array of such utility levels and returns a number. As a result, the state of affair A is better than B if and only if the image of A under this function is higher than the image of B. Totalism and Averagism are paradigms of this type of axiologies; however, these views are infamously known for bringing about undesirable conclusions, e.g. the so-called Repugnant or Sadististic one[12]. Therefore, besides the numbers of people and their well-being levels, it is plausible that other factors, like the distribution of these levels across the population or how specific individuals are affected in each outcome, are morally significant as well and thus need to be considered. To be sure, factors such as the agent's well-being level should not matter; nonetheless, we shall see how the well-being of certain groups of people, what is called "personal goodness" (Bader, a), could be considered as relevant[13].

To conclude, producing an axiology function that considers every relevant aspects as well as that avoids undesirable implications is extremely difficult, if not utterly impossible. Given that we lack a complete description of a feasible axiology function, even if we exactly knew the set of maximally relevant factors, the only viable option would be that of making pairwise comparisons of similarly enough states using dominance principles. The use of dominance principles when comparing two objects represented by multi-dimensional arrays is an effective procedure; nevertheless, selecting the appropriate dominance principle for each comparison of a state vis-à-vis another is a delicate and often tricky matter, as the following

---

[12]For a survey of the most influential population axiologies and respective objections see Greaves (2017)

[13]Temkin refers to the idea of taking an "agent-neutral" or "impartial perspective" when discussing morally relevant factors; this stance follows naturally from adopting an impersonal view. In §5 we shall see that it is possible to commit oneself to neutrality or impartiality but at the same time to caring about personal goodness

example shows.

Multiple factors at play, a unique relation obtains: greater volume

Imagine a scenario where there are different objects that we want to compare wrt their volume. Volume is a gradable physical property that all three-dimensional objects instantiate; as such, "having a greater volume" or equivalently "being more voluminous" could be regarded as a binary relation associated to a multidimensional concept. Ideally, our goal is to determine the object(s) having the greatest volume. Unfortunately, our set of objects is very heterogenous. Some objects are prisms having the same height, but different bases. A subset of them has a squared-base, whereas others have as base other regular polygons. Another subset comprises irregular-shaped objects; this is in turn a subset of those objects of which we know the material out of which they are made. There is only a small number of different materials, so there are at least two objects made of the same substance. The only measuring instruments at our disposal are a balance scale and a fixed number of pairs of identical, crystal-clear, jars of water. Suppose, for the sake of the example, that we are completely unaware of the standard formula to work out the volume of prisms, nor that density is the ratio between mass and volume.

Lacking a function that returns the volume, we cannot but precede by making pairwise comparisons. The only moves available to us are: (I) comparing the base edges of two prisms vis-à-vis each other; (II) comparing the dishes of the balance scale and (III) comparing the amount of water displaced in each jar when an object is immersed. Suppose that this very last move is available only for a limited number of objects, those that are water-proof. From performing (I-III) we can put forth the following dominance principles, which we think are intuitively justified:

  (i) If A and B are prisms with a square base, if A's base edge is longer then B's, then A is more voluminous than B

  (ii) If A and B are made of the same material, if A is heavier than B, then A is more voluminous than B

  (iii) If A and B are water-proof, if the amount of water displaced by immersing A is greater than that displaced by B, then A is more voluminous than B

Suppose that there is no pair of objects that satisfies the antecedents of (i), (ii), and (iii) and no pair that does not satisfy at least one of them. It is clear that the factors relevant for the comparison change according to the type of objects being compared. For the class of prisms, the volume-tracking feature is the

length of the base edges, a property that irregular-shaped objects clearly lack[14]. Similarly, whenever we are comparing two objects made out of the same material, a key factor is their weight, which an immaterial prism would not instantiate. A consequence is that, although volume is a property that can supervene on different sets of lower-order properties of objects and whether the "more-voluminous-than" relation obtains or not depends on different factors according to the class of the relata, the "more-voluminous-than" relation that emerges from the principles above is asymmetric[15].

What about transitivity? I think it is pretty obvious that "being more voluminous than" is transitive. Yet, for all that we can know from (I)-(III), it could be that many direct comparisons are not possible: if A and B are both squared-based prisms, B and C are made of the same material, which is different from A's, and C is irregularly shaped and is not water-proof, then there is no principle among (i)-(iii) that allows us to compare A and C vis-à-vis each other. However, if we remind ourselves that the purpose of dominance principles is that of deriving a comparative claim just by looking at differences in some volume-tracking properties instantiated by objects having similar supervenience bases, then the mere fact that there's no dominance principle that applies should not entail that the "more-voluminous-than" relation does not either. Indeed, dominance principles are sufficient conditions, but not necessary ones, for a comparative relation to hold.

Moreover, it is very easy to get dominance principles wrong, even when just a single dimension is involved. For example, suppose that instead of (i) we had:

(i*) If A and B are prisms, if A's base edge is longer then B's, then A is more voluminous than B

(i*) seems prima facie plausible, but, in fact, turns out to be wrong. If the bases are different polygons, say a square and a hexagon, then it is not longer true that a longer edge is a sufficient condition for a greater volume of the prism[16]. However, that (i*) is false might become evident only if the "supervenience recipe", viz. the formula, to derive the surface of a generic polygon given the side is known; or else if (i*) is judged less plausible than other dominance principles that considered together with it form an inconsistent set, for example by violating asymmetry, which we have seen to be the weakest condition for a "more-F than" relation. Indeed, if the difference in the length of the base edges is substantial, (i*) is true: a square whose edge is ten times that of an hexagon has surely a greater surface. But since we take

---

[14]Indeed, if we imagine these objects as nonpolyhedrons then there is no base-edge factor whatsoever.

[15]Even though it is not possible to logically derive asymmetry, I think it's enough to note that each principle involves a relation that is asymmetric (i.e. "heavier than", "longer than", "greater amount of displaced water") and that for asymmetry to fail one would have to present a plausible counterexample

[16]The ratio between area of base and the length of the edge changes according to the particular shape, affecting the volume as well, which is a function of the area of base; for example, a square with edge 1.25 m long has a smaller surface area than a hexagon with an edge 1 m long

the principle to be applicable to any pair of prisms regardless of the ratio between the respective base edges, even a small difference along the relevant dimension must be accounted for.

This example has illustrated that a comparative relation that takes as sufficient conditions distinct dominance principles that involves different factors is not in itself problematic for transitivity; as long as the dominance principles correctly track how a change to the set of volume-making features affects the higher-order property volume, there is no threat of asymmetry or of nontransitivity at all. If, on the other hand, we overlook some factors relevant to the volume of an object in a specific case, e.g. the relationship between the edge and the surface area in prisms, then it is very likely that dominance principle formulated for that dimension is a mistaken one, leading to cycles or violations of transitivity.

3   Against transitivity: the Spectrum Argument

The so-called Spectrum Arguments à la Temkin-Rachels purport to question the the transitivity of the better-than relation by presenting the reader with a very persuasive counterexample. Arguably, an argument analogous to that for morally betterness could also be run for more-F/F-er than relations concerning a nonnormative concept[17]; all that is required is a multidimensional adjective. In the light of the previous sections regarding the meaning of better than and the discussion of multiple-factor properties, it is time to turn to a locus classicus where such theoretical considerations are put at work, revealing a serious threat to transitivity. In this section I shall provide a condensed analysis of the Spectrum Argument in its most general form, then address some of the most influential responses elicited by it.

Some preliminaries

First of all, a Spectrum Argument could be mainly run for concepts that come in degrees, so that we can define a spectrum of different levels at which the relevant property can be instantiated. For example, it might not be impossible, but it is surely less natural, to generate a spectrum for concepts like "dead" or "true". On the contrary, properties like "cold" or "being a democracy" generate a spectrum of degrees of instantiation, therefore it is meaningful to claim that an object is colder than another and that a government is a democracy to some extent. Secondly, the underlying concept is a multi-dimensional

---

[17]For example, Temkin (2012, §9.2.3) runs a spectrum argument for the relation "hairier than"; see also Wasserman and Hare (2017)

one[18], so that it is possible to move along the spectrum just by acting on a single dimension[19]. "Being healthy" is a higher-order property that depends on several factors; although a mere 5-point difference in one's triglycerides level does not necessarily make a person less healthy than at the time when she had the previous blood test, a 50-point difference surely can. Thirdly, it is possible to remain on the same point of the spectrum by acting on more than one dimension at the same time[20]. Two pianists could be talented alike, one having a great technique but an average interpretative ability, whereas for the other the converse happens to be the case.

If there are concepts that satisfy all these three conditions, I think it's pretty uncontroversial that moral goodness is on this list; the minimal, initial requirements to devise a Spectrum Argument are easily met.

The inner workings of a Spectrum Argument

Besides the above assumptions, the Spectrum Arguments envisaged by Temkin (1996, 2012) and Rachels (1998) require a set of further premises that specifically apply to the case where the concept at hand is that of goodness[21], so that by comparing two objects we are asking which one is better[22].

The simplest type of Spectrum Argument uses only two dimensions, say pain intensity and duration of the experience, or individual well-being level and number of people in a population. Thus, more formally we can identify each object of comparison, specifically a state of affairs, by a pair (α, β), viz. a pair of points on the two spectra of generated by each dimension. Indeed, more than the type of factors involved, what is crucial to develop the argument is that each factor can form a discrete spectrum in turn, meaning that it is possible to move from a point of the spectrum to another by acting finitely-many times on the initial degree that the factor dispalyed. For example, if a relevant dimension is the number of people, we could increase or reduce that number; similarly, we can intensify or alleviate a pain or a pleasure.

Secondly, one needs to assume the existence of two points α_0 and α_n (resp. β_0 and β_n) that are pretty far from each other along the spectrum, yet connected by a chain of $n \in \mathbb{N}$ changes, and which are such that,

---

[18]Properties that grounds an instantiation of the concept don't have to be spectrum-generation apt, they could be an all-or-nothing matter; however, for simplicity, in what follow I shall assume that each of these lower-order properties is also a gradable one.

[19]If, no matter how much a factor varies, the property F gets instantiated to the same extent, then that factor is not relevant for the ascription of F. On the other hand, the fact that *any* change, however small, along an arbitrary dimension results in a change along the spectrum it's a stronger condition to require, which needs to be discuss within a specific context as defined by the nature of the property itself

[20]That is, a change to one factor can offset a difference in another one; as above, no further assumptions regarding under which specific conditions such a such trade-off is possible, are made

[21]To be understood in a moral or prudential sense

[22]For the sake of simplcity, I shall be assuming that we are always interested in the *goodness* level of each object, even though, more properly speaking, some versions of the spectrum argument, like the famous "Hangnail for Torture", deal with a range of *badness* levels, rather than goodness ones. For the same reason, in what follows I shall avoid using the relation "worse than", regarding it just as the dual of "better than"

if the degree of the other dimension is kept constant, there is a substantive difference in their respective goodness levels (e.g. the goodness of (α_n, β) exceeds by far that of (α_0, β)). For instance, we could move from 1 month to 60 years of the same-level pleasurable experience, or range from a life of muzak and potatoes to one of Mozart and Sachertorte led by the same number of people, by making a finite number of well-being-related improvements. The goodness levels of the initial and the final states of affairs are manifestly different.

As a result, given a state of affairs (α_0, β_0), it is possible to start a chain of further $n \in \mathbb{N}$ states of affairs (α_i, β_j), with $i, j \in \{1, ..., n\}$.

Define such a chain:

α-**ascending** if it such that for all $i \in \{1, ..., n\}$, for a fixed j, (α_i, β_j) is better than (α_{i-1}, β_j)

α-**descending** if it such that for all $i \in \{1, ..., n\}$, for a fixed j, (α_{i-1}, β_j) is better than (α_i, β_j)

Similarly for the β-dimension. Moreover, define a chain starting at (α_0, β_0) and ending at (α_n, β_n) "ascending (simpliciter)" iff:

(i) It is either α-descending and β-ascending or α-ascending and β-descending

(ii) (α_i, β_j) is better than $(\alpha_{i-1}, \beta_{j-1})$ for all $i, j \in \{1, ..., n\}$,

Call it "descending (simpliciter)" if

(i) It is either α-descending and β-ascending or α-ascending and β-descending

(ii) $(\alpha_{i-1}, \beta_{j-1})$ is better than (α_i, β_j) for all $i, j \in \{1, ..., n\}$,

Intuitively, the existence of such chains means that a small change along one dimension can trump one made to the other one, yet in the opposite direction.

The third and last claim to make is made of two conjuncts: the first states that there exists either an ascending or a descending chain of states of affairs. The second says that the chain is such that, if it is ascending, then the initial point (α_0, β_0) is better than the final point (α_n, β_n); if it is descending, (α_n, β_n) is better than the initial point (α_0, β_0).

If the better-than relation were transitive, then the starting-point of a descending chain would be better than the final one, and the end-point of an ascending chain would be better than its origin. Therefore, claiming that the relation "being better than" is transitive is inconsistent with the set of assumptions made by the Spectrum Argument. On the pain of having to deal with an inconsistent concept, at least one of these claims must be rejected.

Consequences of the Spectrum Argument

The first thing to notice is that for the Spectrum Argument to be a real challenge to transitivity, one must provide a description of a feasible thought experiment involving states of affairs determined by factors that, at least prima facie, satisfy the requirements introduced above. Only once the details have been filled in, the plausibility of each claim, on its own or considered together with the others, can be discussed and evaluated. Transitivity is a universal property, in the sense that either a relation is always transitive or it is nontransitive, regardless of the specific context or set. Hence, in order to achieve the result that better than is nontransitive, it is enough to find a single concrete counterexample where the transitivity axiom is the least plausible one. On the contrary, the other claims concerns existential properties: the fact that one does not find intuitive that there is a descending chain or that we can range from a point of the spectrum to another taking a finite number of steps in that given example does not in itself disprove that such features cannot be instantiated whatsoever. However, the burden of coming up with a powerful-enough example of the Spectrum Argument is on arguably the opponents of transitivity. In the rest of this section I will briefly survey some of the most common objections to the existential claims, whereas a positive case for transitivity-failure will be addressed in the next section.

Some responses in defense of transitivity

Most of the instances of the Spectrum Argument describe changes of very different magnitude along each dimension: the intensity of pain in $(\alpha_{i-1}, \beta_{j-1})$ is slight less severe than in (α_i, β_j), yet it is meant to last for a number of years which has doubled since β_{j-1}. Voorhoeve (2008) argues that, as a matter of empirical evidence, people tend to focus more on the dimension that varied the most, hence shifting the weight assigned to each factor as we move along the chain of states of affairs. It is no surprise that when confronting a choice between opposite end-points, different psychological factors come into play, and the common intuition is to make a judgment which is the opposite of that suggested by the pattern of the chain. Similarly, one could argue that our intuitions regarding very large numbers or uncommon experiences, like excruciating torture, are not reliable (on this see, among others, Broome (2004, 2003, p.54-49)). Therefore, since we cannot effectively imagine what would be like to endure two years of torture or 60 years of a hangnail, there is no sensible, intuitive judgment that could tell us which one is better. Granting that it is meaningful to trust our intuitions when dealing with scenarios of this kind, it is possible to reject the first claim by arguing that, say, eating a piece of cake and listening to Mozart do not merely vary in pleasure intensity, but are pleasurable experiences of an utterly different kind, and as such

they are not points of the pleasure-spectrum connectable thanks to a step-by-step approach[23]. Moreover, instead of having gaps just in the spectrum of a relevant dimension, one could suppose the occurrence of discontinuities at the goodness level as well. It could be the case that the two points of the goodness spectrum are so qualitatively different that there is no appropriate change in the other dimension that can offset the change occurred in the other one, e.g. no amount of people with a life barely worth-living can counterbalance a population of few blessed ones. Handfield and Rabinowicz (2018) define this phenomenon "radical inferiority": according to the Spectrum Argument, (α_0, β_0) is radically inferior/superior to (α_n, β_n). On the contrary, adjacent states like $(\alpha_{i-1}, \beta_{j-1}), (\alpha_i, \beta_j)$ must be "exchangeable" in order to produce a descending/ascending chain. In fact, if we can find a pair such that the two adjacent states are "radically incommensurable"[24], then the third claim does not hold anymore.

Clearly, to accept that two adjacent states of affairs could be radically incommensurable, one needs to account for the existence of a meaningful threshold that has allegedly been crossed. That said, it could very well be the case that the position of such a barrier is vaguely defined. Indeed, a hint of vagueness seems to pervasively characterize Spectrum Arguments. On one hand, when adjacent states are described, phrases like "a slightly less-intense pain", "a slightly less-serious illness" which definitely involve vague terms, are employed. Even more importantly, the betterness relation itself seems to be affected by vagueness, so that there exist borderline cases where it is not clear which state is better than the other[25].

Vagueness in predicates often results in Sorites sensitivity. A common objection to Spectrum Arguments is that they are a form of the Sorites Argument, even if in disguise. The problem with the Sorites-charge is not that, compared to standard Sorites Arguments, more dimensions are involved, but that each state is predicated to be different from the previous one, for either of the two adjacent states is better than the other. However, if the predicate "good" can be easily regarded as sorital, then making a minor change to some of the good-tracking factors won't make an outcome "not good" all at once; yet, this is the case until we reach a point where it would be highly counterintuitive to judge that state of affairs as "good". Thus, one option advanced by Nebel (2018), is to frame the Spectrum Argument in Sorital terms by using the fact that, arguably, some states of affairs of the Spectrum are ascribable the property of being good or bad simpliciter. On the other hand, Pummer (2018) proposes to recast the required Spectrum-argument

---

[23]This was Griffin's view as reported in Qizilbash (2005)

[24]The meaning of this phrase is that along one dimension, say α, there is a certain threshold t such that when α_{i-1} is below/above t, there is a value β_j such that the pairs $(\alpha_{i-1}, \beta_{j-1}), (\alpha_i, \beta_j)$ are incommensurable, i.e. neither is better than the other nor their goodness levels are equal

[25]Note that incommensurability in comparisons is different from the indeterminacy due to vagueness: if the former might be outcome by a sufficient improvement along some dimension, that the latter could disappear thanks to such a move remains doubtful (Handfield and Rabinowicz, 2018, p.2384). More specifically, Handfield (2014) regards the latter to be a special case of the former.

assumptions so as to avoid the reference to transitivity and replace it with a claim that is plainly Sorital. Lastly, according to Silk (2019), evaluational adjectives, including moral ones, show sorites-susceptibility even in the comparison form. Thus each step along the chain is justified by the the claim that a small difference along some dimensions, if appropriately compensated, can result in a sequence of either better or worse states, whose level of goodness is not relevant. Hence, if we assume that there's a threshold, which could be reached by adding up such small changes, beyond which it becomes undefined whether a state of affairs is better than another, then by transitivity of better than we could derive a betterness relation between the two end-points of chain; such a relation is, if not absurd, at least counterintuitive to hold.

To sum up, there have been several responses to Spectrum Arguments that aim at showing the implausibility of one of the three claims, especially where analyzed in a wider context. Nonetheless, none has regarded the transitivity claim as the most implausible one; in the next section we shall see some of the positive considerations adduced to explain and justify a failure of the transitive property for the better-than relation.

4   In favor of nontransitivity: axiological relativity

Besides merely challenging transitivity by constructing valid arguments based on a set of inconsistent claims, Temkin (2012) devotes several chapters to making a diagnosis and an etiology of the intransitivity of better than. Essentially, the idea is that if the factors relevant to determine whether a state of affairs is, all things considered, better than another changes according to each comparison-pair, then it should be no surprise that the better-than relation is nontransitive: from aRb and bRc we cannot infer aRc simply because $R_{\{a,b\}}$ might involve some R-tracking factors that are not applicable when we compare b with c or a and c. To be sure, several binary relations, for example "being the neighbor of" for people or the "beating" relation on a set of sport teams are nontransitive because the factors relevant for establishing, say aRb are not at play when relating a and c. Nonetheless, neither of these relations involves, conceptually, a level of "neighborness" or of "beatness" that could be ascribed to each relata and compared, in a way as described in §1 for comparative predicates, including "better than". This is key, for the nontransitivity of the latter cannot be convincingly justified by an account that does not mention the goodness of a state of affairs.

In this regard, Temkin puts forth two competing views on the (ethical) value of an object of comparison,

the Internal Aspect View (IAV) and the Essentially Comparative View (ECV). As we shall see, if we adopt the former, then the transitivity of better than follows, whereas endorsing the latter is compatible with nontransitivity. I shall discuss them in turn, focusing more on ECV and also giving some concrete examples.

Two competing views about ethical value

The Internal Aspect View

When we face the task of evaluating the moral value of an outcome, a fairly natural and intuitive view that comes to mind is to assess its several good-making internal features. More specifically, the morally-relevant properties of states of affairs are termed "moral ideals"; I shall postpone a proper discussion of moral ideals until the next section. For the time being, let us consider an outcome A and a set n of relevant moral ideals $\{I^1, ..., I^n\}$ for A. According to IAV, we can think of any state of affairs A as an n-dimensional array $A = \langle I^1(A), ..., I^n(A) \rangle$, wherein each element $I^i(A)$ describes how A fares wrt the morally-relevant dimension I^i. Hence, the all-things-considered goodness of A, $\mathcal{G}(A)$ is a function that from a multi-dimensional argument returns a single value. It is clear that $\mathcal{G}$ depends only on A's score along the n dimensions and how the n scores are combined. All and only A's internal properties matter, since there is no reference to other possible alternatives: the goodness level of an outcome is regarded as an intrinsic property. As a consequence, if we assume that for any set of outcomes $O_1, O_2, ..., O_n$ there is a total order on their images under $\mathcal{G}$, then the transitivity of the betterness relation on the set of states of affairs at hand is a straightforward consequence. Under IAV, the claim that "better than" is nontransitive is not just implausible, but plainly false. The denial of IAV is then a necessary condition for nontransitivity.

The Essentially Comparative View

On the other hand, it could be the case that at least a subset of relevant moral ideals is essentially comparative; this means that these ideals conceptually require at least another alternative to be contrasted with a given outcome A. Therefore, for any such essentially comparative ideal I^i, $I^i(A)$ supervenes not solely on the internal properties of A, but on the properties of some states of affairs in the relevant comparison set $\mathbf{C}$. As a result, the all-things-considered goodness level of A, $\mathcal{G}_{\mathbf{C}}(A)$, depends on the set $\mathbf{C}$ of alternative outcomes that could be compared with A. As this set changes, so could $I_{\mathbf{C}}(A)$ and $\mathcal{G}_{\mathbf{C}}(A)$ in turn.

As a consequence, the choice of **C** becomes axiologically decisive: for example, according to a "pairwise normative theory" (Cusbert, 2017), the relevant **C** is just comprised of the pair of objects being compared. Hence, it is possible that $\mathcal{G}_{\{A,B\}}(A)$ exceeds $\mathcal{G}_{\{A,B\}}(B)$, $\mathcal{G}_{\{B,C\}}(B)$ exceeds $\mathcal{G}_{\{B,C\}}(C)$ and $\mathcal{G}_{\{A,C\}}(C)$ exceeds $\mathcal{G}_{\{A,C\}}(A)$. The possibility of such cycles makes it prima facie plausible that better than is intransitive. However, adopting ECV is not sufficient to reject transitivity; arguably, it is not even a necessary condition. That said, making the goodness level or even the score on each single dimension be a function of a comparison set tremendously complicates the axiological picture. Thus, the move from IAV to ECV needs to be carefully and duly justified and explained[26] .

Some examples of uses of ECV

Although IAV seems more natural, ECV is not wholly unfamiliar. As a matter of fact, whether an object could be ascribed a given property is often the result of comparing it to other relevant objects. To predicate a gradable adjective at the positive form, a so-called "standard of comparison" (Kennedy, 2007) is required. For instance, saying "my car is expensive" means that the speaker deems that the cost of her car is higher than the average cost of a set of cars, with this set being fixed by the context of utterance. If the car is a Ferrari, it is surely more expensive than the average price for a car, period. Yet, it could be that the car is a Fiat 500, which is less expensive than the average but has a higher price compared to cars of a similar size and specs. However, this context sensitivity is not particularly problematic, for whenever the situation might look ambiguous, specifying the relevant context avoid stating inconsistent claims: a speaker can consistently utter both "my car is expensive, for being a city car" and "mine is a city car, not an expensive one".

Similarly, even though the price of a given car could be regarded as an internal feature that does not depend on the comparison set, the expensiveness level could not. To better illustrate the point, consider the way the performance of candidates taking standardized tests, like the GRE, is evaluated. Besides the actual score, viz. how many correct answers the candidate has given, the goodness level is given by a percentile, that is to say a position on a distribution of the scores of a relevant group of test-takers. Thus, it could be the case that, due to a different comparison set, a candidate who scored 165 out of 170 points could be GRE-worse than a candidate who scored 163 in a different year. However, once it is settled that it is the percentile wrt to a certain comparison set that uniquely defines the all-thing-considered

---

[26]Note that both views are compatible with the account of betterness relation given in §1

GRE-goodness level, instead of considering just the GRE score simpliciter[27], then a situation like the above is not in the least paradoxical, nor would it generate cycles or transitivity failures.

Assigning a particular degree by taking into account not merely an object's internal features measured according to a common unit, but also according to their relative value, is a common strategy, from essay-grading in schools to choosing among job applicants. However, note that in these cases the relevant essentially-comparative ideal is not essentially *pairwise* comparative, which would force one to directly compare two outcomes as the only way to establish their relative ranking (Temkin, 1997, p.304). Instead, even when the comparison is between an object vis-à-vis another, the relative goodness of each depend on all the other alternatives. This is crucial, for it ensures that the better-than relation is complete on the relevant comparison set, so that we can always make pairwise comparisons and establish whether an object is better than another. Moreover, for a given comparison set $\mathbf{C}$, for any outcome A there's a unique $\mathcal{G}_{\mathbf{C}}(A)$, so that we can get a total order on the set $\bigcup_{A \in \mathbf{C}} \mathcal{G}_{\mathbf{C}}(A)$. To be sure, $\mathcal{G}_{\mathbf{C}}$ could be a complex function that depends on several factors. For instance, in soccer tournaments the relative goodness level of two given teams A and B is based on several extrinsic factors, including, but not limited to, whether or not team A bit team B. As the relevant $\mathbf{C}$ is the set of all teams taking part to the tournament, how A and B fared against other teams has to be considered as well. Thus, a pairwise comparison limited to A-vs-B games would inevitably disregard some meaningful factors, and any relation of betterness obtaining would not be an all-things-considered one.

To conclude, it is worth noticing that the choice of $\mathbf{C}$ depends also on the set of the relevant good-tracking factors: in the simplest cases, for an arbitrary object $A \in \mathbf{C}$ there is a unique function $\mathcal{G}_{\mathbf{C}}(A) = f(\langle I_{\mathbf{C}}^1(A), ..., I_{\mathbf{C}}^n(A) \rangle)$. In this case, $\mathbf{C}$ would be homogeneous, meaning that same types of properties constitute the goodness-supervenience base for each object in $\mathbf{C}$[28]. The axiological challenge specifically presented by the adoption of ECV occurs whenever the objects to be compared form a heterogenous collection; in that case, an all-encompassing comparison set is not a viable option, nor is giving a unique description of the goodness function for all alternatives.

---

[27]Arguably, one could give a sort of lexical priority to percentiles, so that numeric scores become relevant only if the percentiles are the same

[28]One might argue that once fixed $\mathbf{C}$, then $\mathcal{G}_{\mathbf{C}}$ could be considered an internal feature of an outcome, just like saying that "Alice being 6-feet tall" is an internal property of Alice; however, strictly speaking her *height* is an internal property, whereas its value is due to a relation to the external object "foot".

Different factors, ergo different relations?

As already hinted above, the threat to transitivity posed by ECV springs from those situations wherein the relative goodness level of an outcome is considered to be meaningful even if it depends on the smallest choice of **C**, viz. a pairwise comparison. Since our focus is a binary relation and not a complete ranking of the options, it seems that some properties (or the lack of) of betterness can emerge from such tête-à-têtês. Spectrum Arguments hinge on the fact that our intuitions might change in each pairwise comparison: the relative goodness of 60 years of hangnail exceeds that of two years of excruciating torture, yet not that of a slightly more intense pain endured for 40 years, being the goodness factors that come into play in the former comparison different from those that are relevant for the latter. A first response to this is that of providing a fine-grained analysis of the comparison sets or equivalently of the betterness relations concerned[29]. On this view, in any cycle as the followings:

$$\mathcal{G}(A) >_{\{A,B\}} \mathcal{G}(B), \qquad \mathcal{G}(B) >_{\{B,C\}} \mathcal{G}(C), \qquad \mathcal{G}(C) >_{\{A,C\}} \mathcal{G}(A)$$

a different better-than relation ($>_\mathbf{C}$) would obtain for each comparison, so that the possibility of such a cycle does not constitute a counterexample to the transitivity of *genuinely* better than. Indeed, all that the above cycle might reveal is that there is no object that is all-things-considered better than any other in the comparison set.

For instance, let us consider the Rock-Paper-Scissors game. According to the rules, the "better-to-play-than" relation ($>$) produces the cycle below:

$$R >_{\{R,S\}} S, \qquad S >_{\{S,P\}} P, \qquad P_{\{R,P\}} > R$$

Nevertheless, relativized to the set of all alternatives $\{S, R, P\}$ that one can play, all options are equally good, there's none that would be a better choice than another. Hence, it seems clear that the three occurrences of $>$ above are not instances of all-things-considered betterness as described in §1 and §2.

I believe that fine-grained-style replies are on the right track to dispel worries related to apparent violations of transitivity, yet further considerations are in order.

On one hand, we can regard the choice of the relevant comparison set as a function of the choice of those factors that are significant for making that very comparison. As a result, given that when facing a

---

[29]This is the move advanced, among others, by Voorhoeve (2013) and Handfield (2016)

pairwise comparison the set of relevant moral ideals $I_{\{A,B\}}$ is not identical to the sets $I_{\{B,C\}}$ or $I_{\{A,C\}}$[30], one might conclude that what prima facie looks like three tokens of the same type of relation, under further reflection the three reveal to be instances of conceptually distinct relations. For example, consider the following sentences:

(i) David is fairer than Solomon[31]

(ii) Solomon is fairer than the US government

(iii) ??David is fairer than the US government

The twofold meaning of the polysemy adjective "fair" can be explained by the fact that the fairness-tracking properties in the (i) and (ii) respectively form two distinct, possibly even disjoint, sets. In (i), fairer is a synonym for "more handsome", whereas in (ii) for "more just"; obviously, different factors are significant to compare physical appearances and impartial policies. It is hard to say that there should be an expectation that (iii) holds, for otherwise we would be facing a counterexample to the transitivity of "fairer than". In cases like this, despite the fact the same word figures in both sentences, the truth-making conditions of (i) and (ii) are so different in kind that it is clear that we are dealing with two ontologically different relations. On the other hand, this is hardly a universal conclusion.

Recall the greater-volume example of section §2. In that case, the relation "more voluminous than" was grounded by a different set of properties, according to the sort of objects being compared. Weight and base-edge length are factors of a very different nature indeed; nevertheless, we were not inclined to distinguish the more-voulouminous-than relation between prisms from that holding between objects of the same material. The same property, vo

To conclude, it is very difficult to exactly pinpoint the identity conditions for a relation, and this is not a topic that could be even merely sketched here. Yet, I think it is fairly safe to say that it is not a necessary condition that the property-tracking factors constitute the same set for every pair for which the given relation is supposed to obtain. Therefore, those who invoke a fine-grained move to explain away a failure of transitivity cannot limit themselves to note that since different moral ideals come into play, it follows that the relations obtaining are ontologically different in kind. If, as in the above example with "fairer", the argument displays an equivocation, this must be accounted for not just by a difference in

---

[30]It could also be the case that the sets of factors are the same, but that different weight is given to each, viz. even if the same outcome B figures in both comparisons, there are two distinct goodness functions: $f_{\{A,B\}}(\langle I^1_{\{\mathbf{A},\mathbf{B}\}}(B),...,I^n_{\{\mathbf{A},\mathbf{B}\}}(B)\rangle)$ and $f_{\{B,C\}}(\langle I^1_{\{\mathbf{B},\mathbf{C}\}}(B),...,I^n_{\{\mathbf{B},\mathbf{C}\}}(B)\rangle)$

[31]Apparently, David is considered one of the most handsome characters in the Old Testament

the supervenience base, but by the fact that this base grounds an altogether different property. This is much more difficult to show; whether there are plausible ways to adopt ECV and make sure that the same all-things-considered goodness is the only property being evaluated in each comparison will be discussed in the next section.

5   Goodness: A personal matter

Dominance principles and value

Recall that according to an ECV of ethical value, the relative goodness level of a state of affairs depends also on how the state fares relative to another one along those moral ideals that are essentially comparative. More specifically, *pairwise* essentially comparative ideals allow us to compare two states vis-à-vis each other and establish which is better wrt to that given factor. Since the significant factors are several even when comparing just a pair of outcomes, and given that we often lack a definite recipe to combine such good-making factors, pairwise essentially comparative ideals work best as dominance principles, namely as principles that apply only in those cases where the two states score the same along other relevant dimensions. Examples of this type of ideals are Pareto Principle, Rawlsian Maximin and even a principle according to which an outcome is better than another, utility-wise, if the well-being of already existing people has increased[32].

Compared to other ideals like equality or total utility, the purpose of the above principles is to provide a relative ranking of two or more states of affairs in a way that is meant to track our intuitions regarding moral goodness. If one believes that it is important to improve the conditions of those who are worse off, then it is uncontroversial that, ceteris paribus, a state of affairs where a higher level of well-being accrues to this group is all-things-considered better than another wherein their well-being level is lower. This sort of principle retains the same intuitive flavor as a the idea that a heavier object made out of the same material is more voluminous than one that is lighter.

Dominance principles are of great importance in all those cases where one lacks a clear supervenience recipe, or well-defined mathematical formula that outputs the extent by which a gradable property is instantiated from a set of factors. On the other hand, more often than not, the scope of dominance principles ends up being a very narrow one. In order to minimize the the risk of misjudging the relative all-things-considered goodness by appealing just to one small difference, the comparison set must be

---

[32]See Temkin (2012, Ch. 12) for a detailed discussion of essentially comparative ideals

comprised of outcomes that fare very similar, if not identical, in all other respects. Consequently, it could happen that, within a set of more than two alternatives, only some dominance principles apply and only for some subsets of such alternatives. That said, this is not in itself a direct cause of intransitivity, as the volume example shows.

Person-Affecting Principles

At the root of essentially comparative ideals is the view that how different groups of people are affected across alternatives is a morally significant factor, and that such comparative well-being cannot be captured by IAV alone. First of all, one can distinguish between people who would exist no matter what outcome will be brought about, namely those groups who figure in every state of affairs of the relevant comparison set, and people whose existence depends upon the actualization of a specific outcome. Hence, granted that brining someone into existence does not benefit that person, nor failing to do so would harm her[33], the Narrow Person-Affecting View (PAV) is the view that

> In assessing possible outcomes, one should (1) focus on the status of independently existing people, with the aim of wanting them to be as well off as possible, and (2) ignore the status of dependently existing people, except that one wants to avoid harming them as much as possible.[34]

As noted by some commentators[35], since a necessary condition to harm a group of people is failing to maximize their well-being and that nonexistence in an outcome doesn't count as a harm, Temkin's definition of Narrow PAV amounts to the claim that if there is an alternative O_2 that would be better for a given group compared to O_1, then that group would be harmed in O_1 and this fact affects the goodness level of the outcome itself. In short, the better alternative is the one that minimizes comparative harm. When we compare populations, it could happen that we have only partial group-overlappings, i.e. some people exist in some states of affairs but not in others, so that an outcome could be better than another just in virtue of being better *for* a specific group; as a result, it is very easy to generate prima facie instances of nontranistivity for betterness or cycles. The most cited example is the one showed in Figure 1, originally owed to Parfit[36].

---

[33]This claim is often referred to as the "neutrality intuition": the mere fact that in a given outcome a person exists or not is regarded as neutral, as a matter of personal goodness.

[34]Temkin (2012, p.417)

[35]Among others, Roberts (2014) and Ross (2015)

[36]As reported by Temkin (2012, p.428-429). In Parfit's original example the pair of people involved were supposed to have a normal life (the highest column) and an handicapped one respectively, so as to rule out a possible evolution in time from a state to another; the picture below is from Ross (2015, p.435)
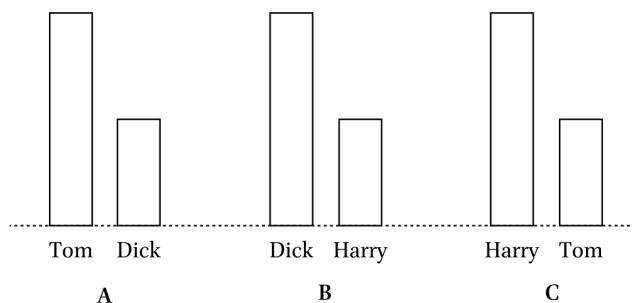
Figure 1: Tom, Dick and Harry

According to a Narrow PAV, Tom is an independently existing person who fares better in A than in C; thus, since all the other factors are equal (Dick's or Harry's well-being would be equally maximized in either A or C), A is all-things-considered better than C. Similarly, C is better than B, and B is better than A. The result is a cycle generated by a failure of transitivity of all-things-considered betterness. Temkin (2012, p.431-432) himself notes that on Narrow PAV there is a straightforward move available to avoid producing such a cycle, for it is clear that when the three alternatives are considered at once, anyone among Harry, Tom and Dick would be harmed to same extent regardless of which alternative is chosen. Therefore, being the amount of comparative harm the same, by definition the outcomes are all equally good.

It is worth noting that even though we can avoid a cycle, the betterness relation would still be nontransitive, in the sense that from the first two claims ($A > C$, $C > B$), one would be expecting that $A > B$ and be disappointed. As discussed in the previous section, such an expectation is a legitimate one only if, although different factors come into play at each comparison, the relation is conceptually the same, for it deals with the same type of goodnesses, regardless of whether the relata are A and B or A and C. Should it turn out that we are comparing the edge base of a square with that of a pentagon, then the the dominance principle employed is not tracking the same comparative property as that tracked by comparing the amount of water displaced. As a result, it could be the case that the widespread failure of transitivity due to endorsing Narrow PAV could be avoided by paying greater heed to the good-making properties that a person-affecting intuition is meant to capture.

Revising PAV to bypass intransivity

It is worth investigating whether Narrow PAV is successful in comparing the moral goodness of two outcomes wrt person-affecting concerns, or whether it needs to be amended and revised, regardless of

transitivity-related issues. I shall conclude this section by briefly discussing versions of PAV that purport to track good-making factors more properly, thus avoiding some charges of intransitivity that Narrow PAV seems doomed to face.

First of all, let us go back to Figure 1. Personally, I see no strong reason to judge A better than C; I think that A, B and C should be regarded as equally good not just when considered at once, but in pairwise comparisons as well. Indeed, if there weren't any specification of the identity of the persons involved in each alternative, no one would have dreamed of judging an option better than the other: A, B and C would be equally good also compared vis-à-vis each other. Therefore, the question to ask is whether identity-related factors should be considered as morally relevant: it's not enough that there's someone who is better off in an alternative outcome, but this person must be the same in both outcomes. Although personal good is arguably an axiologically-relevant dimension, from "the point of view of the universe" the specific identity of the persons to whom the goodness accrues is not. Assumed that there's no evolution from an outcome to another, hence no loss of well-being, and that they are all possible alternatives and none is actual, then there is no significant reason to countenance people's identity in the supervenience base of an outcome's goodness, meant to be assessed from an agent-neutral perspective. Hence, the idea is to revise PAV so as to make it "impartial": following Bader (b, Ch.3) it is possible to put forth a *Weak*[37] PAV that allows to compare permutations of groups of people existing in two outcomes. Since there is a permutation that maps Tom to Harry and Dick to Tom, when A and C are compared neither would be better, on Weak PAV. By doing so, transitivity is not violated.

Similarly, among others, authors like Horta (2014) and Ross (2015) propose refined versions of person-affecting principles that purport to better accommodate some of our moral intuitions, for example regarding the Mere Addition Paradox, with the byproduct of retaining the transitivity of all-things-considered better than. Therefore, the intuitive purchase of a person-affecting view, once properly framed, does not necessarily result in a nontransitive betterness relation, especially considering that in order to make an all-things-considered judgment, person-affecting considerations need to be weighed against other morally significant factors in an appropriate way[38].

---

[37]Sometimes this view is also known as "wide" PAV, to contrast it with the Narrow one

[38]For example see Dancy (2005) on this point

6 Conclusion

To wrap up, the points made in this paper are the followings. Whether the all-things-considered better-than relation between two states of affairs obtains or not depends on the (moral) goodness level instantiated by each state. On a pluralist view, such a goodness is a multiple-realizable property, in this sense: it can supervene upon different set of factors, and according to a different supervenience-recipe, according to the nature of the state of affairs being evaluated. Moreover, the set of significant factors and/or the supervenience-recipe could change according to the relevant set of comparison for that outcome in a given context. This means that there might be no absolute level of the ethical value of a state of affairs, and that we should lean towards a context-relative axiology. This move is especially compelling if great moral weight is ascribed to essentially comparative ideals.

Nonetheless, by allowing for goodness to be relative, there are higher chances of facing at least prima facie instances of nontransitivity of better-than, like those envisaged by Spectrum Arguments. However, this could be explained by the fact that ECV is more prone to equivocate on the type of goodness concerned, thus of being dealing with what are ontologically different betterness relations. If this happens, transitivity might not apply, nor should we expect it so, however.

Sameness of supervenience bases is not a necessary condition for identifying properties and associated comparative relations. Alas, in order to coherently evaluate alternatives, it remains crucial to make sure that outcomes are being compared according to the same property, and this can only be done by taking into account the right sort of good-making factors and they way the combine in each relevant comparison. In this regard, using dominance principles is a convenient strategy to directly compare outcomes that satisfy some similarity conditions in many respects and differ along just a single dimension.

Granted that personal good is a morally significant factor, person-affecting principles should be regarded as dominance principles. This implies that they are very limited in scope and somewhat weak. Admittedly, one of the aims of population axiology is that of refining, checking and possibly expanding the scope of such principles. My sense is that the transitivity of better than is a straightforward consequence of correct principles.

References

R. M. Bader. Person-affecting utilitarianism. In Arrhenius, Bykvist, and Campbell, editors, *Oxford Handbook of Population Ethics*. Oxford University Press, a.

R. M. Bader. Person-affecting population ethics. b.

J. Broome. A comment on Temkin's trade-offs. In Wikler and Murray, editor, *Fairness and Goodness in Health*, Manuscript. World Health Organization, Geneva, 2003.

J. Broome. *Weighing Lives.* Oxford University Press, 2004.

J. Cusbert. Acting on Essentially Comparative Goodness. *Thought*, 6(2):73–83, 2017.

J. Dancy. Essentially Comparative Concepts. *Journal of Ethics and Social Philosophy*, 1(2):1–16, 2005.

C. Dorr, J. M. Nebel, and J. Zuehl. The Case for Comparability. 2020.

H. Greaves. Population axiology. *Philosophy Compass*, 12(11):1–15, 2017. ISSN 17479991. doi: 10.1111/phc3.12442.

T. Handfield. Rational Choice and the Transitivity of Betterness. *Philosophy and Phenomenological Research*, 89(3):584–604, 2014.

T. Handfield. Essentially Comparative Value Does Not Threaten Transitivity. *Thought*, 5(1):3–12, 2016.

T. Handfield and W. Rabinowicz. Incommensurability and vagueness in spectrum arguments: options for saving transitivity of betterness. *Philosophical Studies*, 175(9):2373–2387, 2018. ISSN 15730883. doi: 10.1007/s11098-017-0963-9.

C. Hare. The Great Spectrum Paradox 1. (March):1–27, 2017.

O. Horta. In Defense of the Internal Aspects View: Person-Affecting Reasons, Spectrum Arguments and Inconsistent Intuitions. *Law, Ethics and Philosophy*, 2(2):87–107, 2014.

M. Huemer. Transitivity, Comparative Value, and the Methods of Ethics. *Ethics*, 123(2):318–345, 2013.

C. Kennedy. Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30(1):1–45, 2007.

F. Moltmann. Degree structure as trope structure: A trope-based analysis of positive and comparative adjectives. *Linguistics and Philosophy*, 32(1):51–94, 2009.

J. M. Nebel. The Good, the Bad, and the Transitivity of Better Than*. *Nous*, 52(4):874–899, 2018.

T. Pummer. Spectrum arguments and hypersensitivity. *Philosophical Studies*, 175(7):1729–1744, 2018.

M. Qizilbash. Transitivity and vagueness. *Economics and Philosophy*, 21(1):109–131, 2005.

S. Rachels. Counterexamples to the transitivity of better than. *Australasian Journal of Philosophy*, 76 (1):71–83, mar 1998.

M. A. Roberts. Temkin's essentially comparative view, wrongful life and the mere addition paradox. *Analysis*, 74(2):306–326, 2014.

J. Ross. Rethinking the Person-Affecting Principle. *Journal of Moral Philosophy*, 12(4):428–461, 2015.

A. Silk. Evaluational adjectives. *Philosophy and Phenomenological Research*, Forthcomin:1–35, 2019.

L. S. Temkin. A Continuum Argument for Intransitivity. *Philosophy & Public Affairs*, 25(3):175–210, 1996.

L. S. Temkin. Rethinking the Good. Moral Ideals and the Nature of Practical Reasoning. In J. Dancy, editor, *Reading Parfit*, pages 290–345. Oxford: Blackwell, 1997.

L. S. Temkin. *Rethinking the GoodMoral Ideals and the Nature of Practical Reasoning.* Oxford University Press, jan 2012.

A. Voorhoeve. Heuristics and biases in a purported counter-example to the acyclicity of 'better than'. *Politics, Philosophy and Economics*, 7(3):285–299, 2008.

A. Voorhoeve. Vaulting intuition: Temkin's critique of transitivity. *Economics and Philosophy*, 29(3): 409–423, 2013.

R. Wasserman. Paradoxes of Transitivity *. pages 1–14.