

FEDERICO L.G. FAROLDI\*, ANDRÉS SORIA RUIZ\*\*

## THE SCALE STRUCTURE OF MORAL ADJECTIVES

**SUMMARY:** In this paper we discuss how and whether moral adjectives fit a well-known semantics for gradable adjectives. We first test whether moral adjectives are relative or absolute adjectives. The preliminary results suggest that moral adjectives don't fall neatly under either category. In the second part we tackle the question of the scale of moral adjectives in a more theoretical fashion, i.e. by investigating their possible scales with mathematically precise tools.

**KEYWORDS:** moral adjectives; scalar semantics, multidimensionality, meta-ethics

### 1. INTRODUCTION

“It is morally better to keep a promise than to save a life.” Sentences like these show that moral adjectives such as ‘good’, ‘bad’ ‘(un)ethical’, ‘cruel’, ‘(im)moral’, ‘virtuous’ or ‘despicable’ are gradable, that is, they place their objects on a scale.<sup>1</sup> Can we say more about their semantics, and the types of scales they use? In this paper we apply available tests

---

\* Centre of Logic and Philosophy of Science, Ghent University and Research Foundation – Flanders (FWO). E-mail: federico.faroldi@ugent.be

\*\* Institut Jean Nicod, Département d'études cognitives, ENS, EHESS, PSL Research University, CNRS, Paris France. E-mail: andressoriaruiz@gmail.com

<sup>1</sup> There are thick and thin terms in the previous list. We will not, however, discuss that distinction here (see Väyrynen 2016 for a recent overview of this debate).

to see how and whether moral adjectives fit a well-known semantics for gradable adjectives.

But before that, why focus on moral adjectives? Moral adjectives form an easily recognizable class of words, although this class most likely does not carve a natural joint in language, so to speak. By way of introduction and justification, we note that the present project is inspired by three main trends that are very much alive in the current literature in philosophy and linguistics.

First, recent years have witnessed a surge of literature on the semantics and pragmatics of words that share some (if not many) semantic features with moral vocabulary, such as subjective terms – most eminently, predicates of personal taste, i.e. *fun* and *tasty* (Kölbel 2003; Lasersohn 2005; Stephenson 2007; Stojanovic 2007; Meier and van Wijnbergen-Huitink 2016 *a.m.o.*), and, more recently, aesthetic adjectives too (see Liao and Meskin 2015, Liao *et al* 2016; Stojanovic 2016; Umbach 2016) – modals of various flavors (stemming from Kratzer's seminal contributions, see her 2012) and other purportedly normative expressions, such as *know* or *rational*. To our knowledge, moral expressions have received comparatively little attention.<sup>2</sup>

Secondly, and by contrast, moral vocabulary has long been discussed in meta-ethics. The tradition loosely inaugurated by Moore took very seriously the idea that the inquiry into the linguistic properties of moral words was the right way to gain philosophical insight into the nature of the moral concepts that we deploy. Ordinary language philosophers followed suit, and language-oriented analytic meta-ethics lived something of a heyday that lasted approximately until the second half of the XXth century (Darwall *et al* 1992).

Its decay coincided in time with the emergence of formal semantics, and more specifically, with the appearance of degree semantics, which is the third tradition upon which this project rests. Since the 1970s, much literature in formal semantics has focused on the semantic properties of gradable expressions, often tackling difficult and well-established problems in philosophy, such as vagueness (Cresswell 1976; Klein 1980; Barker 2002; Fara 2003; Kennedy 2007; Lassiter 2016).

---

<sup>2</sup> A curious example is MacFarlane's: his 2014 covers pretty much all the range of expressions that have been discussed in the contextualism-relativism-expressivism debate, but never does he tackle moral adjectives (even though he does devote a whole chapter to the modal *ought*, which is moral under one interpretation).

Thus, our aim is to employ the resources of degree semantics to explore the semantic properties of moral adjectives, starting from the basic observation that these expressions admit comparisons. A worry at the outset, however, is that from the fact that there is a comparative it doesn't follow that there is a scale. If we just consider *good*, for instance, and a general betterness relation, it is somewhat common in the literature to assume that it cannot be a total relation over all alternatives, because of incommensurability or various forms of non-comparability (for an introduction see Chang 1997), while other authors like Temkin (2012) additionally argue that the 'betterness' relation is not transitive. We acknowledge these worries and we hold their discussion, as well as a precise definition of what we mean by a 'scale', until section 3.

We assume, as it's standard, that the semantic value of a gradable adjective  $G$  is a function  $\mu$  from an arbitrary type  $\alpha$  (individuals, for instance) to degrees on a scale:

$$\llbracket G \rrbracket_{\langle \alpha, d \rangle}^{M,w} = \lambda k_{\alpha} \cdot \mu_G(k)$$

The positive form of a gradable adjective  $G_{POZ}$  predicates a property of the degree possessed by  $\alpha$ :

$$\llbracket \text{is } G_{POS} \rrbracket_{\langle \alpha, t \rangle} = \lambda k_{\alpha} \cdot \mu_G(k) \text{ is } P$$

We start (sect. 2) by considering the properties that the positive form of moral adjectives ascribes to their objects. In particular, we test whether moral adjectives are relative or absolute adjectives. To do this, we look at the entailment patterns of moral adjectives, their compatibility with modifiers such as *perfectly* and *slightly* and their sensitivity to comparison classes. The preliminary results point towards the fact that moral adjectives don't fall neatly under either category. Nonetheless, we take the available data to suggest that moral adjectives are relative, even though they display two features that have often been associated with absolute adjectives: they are insensitive to comparison classes and they admit modification by modifiers like *perfectly* and *slightly*. We present a hypothesis to account for these results.

In section 3 we turn to the properties of the measure function  $\mu$ . Thus, we tackle the question of the scale of moral adjectives in a more theoretical fashion, i.e. by investigating their possible scales with mathematically precise tools. We show results of two classes: first,

boundedness properties are compatible both with ratio and interval scale and are, therefore, uninformative; second, we prove that moral adjectives cannot always have a ratio scale because of some non-additive cases. These results pave the way to the multidimensionality hypothesis.

## 2. RELATIVE OR ABSOLUTE?

Gradable adjectives are usually classified as relative (*rel*) and absolute (*abs*), depending on the property that the positive form ascribes to their object's degree (property *P* in the above formalization). The positive form of relative adjectives predicates that the object's degree surpasses a certain threshold, which is an intermediate point on a scale whose precise value is determined contextually; absolute adjectives are those whose positive form predicates that their object's degree is a scale endpoint. *Tall* is a typical example of a relative adjective: to be tall amounts to possessing a degree of height that exceeds a certain threshold, which is an intermediate point on the height scale.

Absolute adjectives are called *maximum* ( $abs_{max}$ ) or *minimum* ( $abs_{min}$ ) standard if the endpoint is the upper or lower scale endpoint, respectively. *Full* is an  $abs_{max}$  adjective: to be full is to have a maximum degree of the property of fullness (namely, to have as much content as capacity); *dirty* is an  $abs_{min}$  adjective, since to be dirty is to possess a non-zero degree of dirtiness.

In addition to this, gradable adjectives are positive or negative, depending on whether modification by *-er* (or *more*) denotes a higher degree on the relevant scale. Thus, *tall* is positive, while *short* is negative; *full* is positive while *empty* is negative; and *dirty* is positive while *clean* is negative.

Interestingly, common tests for this distinction used in the literature do not give stable results when applied to moral adjectives. In particular, while certain entailment patterns suggest that these adjectives are relative, their behavior with respect to PPs denoting comparison classes, as well as their admissibility of certain modifiers suggests that they are absolute (see Liao and Meskin 2015; Liao et al. 2016). Finally, some of these tests give different results for positive and negative adjectives, so we'll consider adjectives of both types, namely

*virtuous/ethical/generous* (positive) and *cruel/despicable* (negative). Let's review these tests briefly.

We'll start by looking at three tests according to which moral adjectives come out as relative (Kennedy 2007, sec 3). First, given that absolute adjectives in the positive form denote endpoints on a scale, the following schemata hold:

if  $x$  is  $abs_{max}$ , then  $x$  could not be  $abs_{max}$ -er.

if  $x$  is not  $abs_{min}$ , then  $x$  does not possess any degree of  $abs_{min}$ -ness.

To test these entailment patterns, we consider whether sentences that do not respect them are acceptable. That is, we consider whether the conjunction of the antecedent clause and the denial of the consequent are coherent. If not, then the pattern holds for that adjective. Consider the following two cases with *full* ( $abs_{max}$ ) and *open* ( $abs_{min}$ ), and compare them with *tall*:

- 1) # The glass is full, but it could be fuller.
- 2) # The door is not open, although it's ajar.
- 3) Natalia is tall, but she could be taller.
- 4) Matheus is not tall, but he possesses some height.

As the examples show, the relevant constructions are not coherent when we use absolute adjectives, thereby showing that these adjectives do not respect those entailment patterns. When we use a relative adjective like *tall* however, those constructions are coherent. Positive and negative moral adjectives pattern like relative adjectives in this test:

- 5) What she did was ethical, although it could have been more ethical.<sup>3</sup>
- 6) Being vegetarian is not virtuous, although it has some degree of virtuousness.
- 7) What he did was despicable, although it could have been more despicable.
- 8) Eating animals is not cruel, although it has some degree of cruelty.

---

<sup>3</sup> Some may balk at this: after all, it seems that, for what she did to be more ethical, she should've done something else altogether. What she in fact did cannot have more or less moral value than it actually has. We share the intuition, but we don't take that intuition to say anything about the scalar properties of *ethical* (or any moral adjective, for that matter). It does however, say something about the objects of moral evaluation.

Secondly, absolute adjectives also satisfy the following entailment pattern:

if  $x$  is more  $abs_{min}$  than  $y$ , then  $x$  is  $abs_{min}$ .

if  $x$  is more  $abs_{max}$  than  $y$ , then  $y$  is not  $abs_{max}$ .

Again, compare for illustration *full* and *dirty* to *tall*:

9) # The shirt is dirtier than the jacket, although the shirt isn't dirty.

10) # The glass is fuller than the vase, although the vase is full.<sup>4</sup>

11) Natalia is taller than Matheus, although she isn't tall / Matheus is tall.

With respect to these two patterns, positive and negative moral adjectives give slightly different results: positive adjectives do not show these patterns, hence they behave like relative adjectives. Negative adjectives show the second pattern, but our judgments about the former are less clear:

12) Although it isn't ethical, animal testing for scientific purposes is more ethical than for cosmetic purposes.

13) Donating to a charity is virtuous, and volunteering for a charity is more virtuous.

14) ?? Although it isn't despicable, hiding your office mate's keys is more despicable than eating their snacks.

Paying bribes is despicable, although accepting them is more despicable.

Using a negative moral adjective like *despicable* in a comparison does seem to suggest that both terms of comparison deserve the positive form. However, we are skeptical that this particular pattern of inference is an entailment. The fact that a sentence like (14), even if marked, is not completely ruled out, suggests that the inference is to some extent cancelable, so it might well arise due to implicature or some other mechanism.<sup>5</sup>

<sup>4</sup> Sometimes constructions like these are acceptable, cf. *the first theater was full, and the second was fuller*. However, when we say such things we are arguably speaking imprecisely. That is, we don't really mean that the first theater is full in the literal sense of not having any free seats, but rather that it is very crowded, or something like that (see Kennedy 2007, p. 23–24).

<sup>5</sup> Bierwisch (1989, p. 89) makes a similar observation about the pair of comparatives *besser (better) / schlechter (worse)*: comparisons with *schlechter*, but not with

The previous observations strongly suggest that moral adjectives are relative. On the other hand, the behavior of moral adjectives with respect to their thresholds' sensitivity to *comparison classes* seems to suggest that they are *not*.

As we mentioned, the thresholds for relative adjectives are determined in context. More specifically, the value of a threshold can be shifted by reference to a comparison class: *tall for a basketball player* and *tall for a 5 year old* establish different thresholds for the positive form of the adjective *tall*. By contrast, absolute adjectives do not show such sensitivity to comparison classes: modification by a comparison class forces the interpretation that the positive form doesn't apply: the PP *for a TV antenna* in *straight for a TV antenna* does not shift the threshold of straightness, but rather is most naturally taken to mean *not straight*.

Liao et al. (2016) observe that *aesthetic* adjectives pattern in this respect like absolute adjectives: rather than shifting the relevant threshold, mentioning a comparison class suggests that the bare positive form doesn't apply (they also note that such constructions – namely 'aesthetic adjective + *for* comparison class' – appear very rarely in *corpora*):

- 16) Anyone who calls someone 'beautiful for an older woman' does not get my love.  
 17) Elegant for a Best Western.<sup>6</sup>

As it's clear, mentioning the comparison class in both cases suggests that the bare positive form does not apply. We think that the same applies to moral adjectives: explicit reference to comparison classes does not shift the threshold but rather suggests that the bare positive form wouldn't apply. Moreover, such constructions sound slightly marked to our ear:

- 18) Giving alms is a generous act (?? for a miser)  
 19) What they did was not despicable (?? for a vile person)

---

*besser*, invite the inference that the positive form *schlecht* (*bad*) applies to both *relata*. However, later (p. 206–207) he says that both comparatives in the pair *schöner* (*prettier*) / *häßlicher* (*uglier*) invite such inference. We acknowledge that the inferences are there, but the fact that they seem in general defeasible suggests to us that they are not entailments, and hence that the hypothesis that these adjectives are relative-standard is not decisively challenged by these observations.

<sup>6</sup>These are Liao et al (2016)'s examples (4a–b), one of the few instances of that construction that they found in *corpora*.

Lastly, the scales of gradable adjectives can be open on either, neither or both ends; and the acceptability of modifiers like *slightly* and *perfectly* (which pick out minimal/maximal endpoints) with positive and negative adjectives has been taken to reveal information about the type of scale lexicalized by a given pair of adjectives (Kennedy 2007). For example, the pair *dirty, clean* lexicalizes an upper-closed scale, as shown by the following pattern:

20) perfectly / ?? slightly clean

21) ?? perfectly / slightly dirty

Interestingly, moral adjectives (in particular, the pair *ethical, unethical*) show the same pattern:

22) perfectly / ?? slightly ethical

23) ?? perfectly / slightly unethical

Kennedy proposes the generalization that adjectives that lexicalize totally open scales are relative, while adjectives that lexicalize scales closed on one or two ends have absolute interpretations. If this generalization is correct, moral adjectives are absolute.

To summarize: Entailment patterns suggest that moral adjectives are relative; while lack of sensitivity to comparison classes and the fact that pairs of moral adjectives admit scale endpoint modifiers suggests that they are absolute.

More tests can and ought to be carried out, but we venture that the hypothesis that moral adjectives are relative is better supported by the data (i.e. the observations about entailment patterns appear particularly telling), whereas tests that suggest that moral adjectives are absolute are based on observations that can be explained alternatively: first, insensitivity to comparison classes *per se* does not show that moral adjectives are absolute, but simply that their thresholds' are rigid in a way that the thresholds of other relative adjectives are not. Why is this so? We may rehearse the following preliminary answer: whether an object falls under a certain moral concept, say *whether an action A is despicable*, depends – in a way to be carefully spelled out – on the moral values of the person who is considering that question. But, crucially, it does not depend on the actions that we may compare *A* with. Thus, it is to be expected that mentioning a comparison class does not change our evaluation.



Secondly, the association of open scales to relative adjectives and of partially closed scales to absolute adjectives is a generalization that has been challenged. In particular, Lassiter and Goodman (2013; 2015) challenge the claim that relative adjectives are associated with an open scale. Moral adjectives may be an exception to the second generalization.

Moreover, the acceptability pattern of modifiers like *perfectly/slightly* can receive an alternative explanation in terms of the *multidimensionality* of moral adjectives. Multidimensional adjectives (Sassoon 2013, 2016; called *evaluative* in Bierwisch 1989) are adjectives that denote properties that can be possessed relative to different respects or dimensions. The pair *healthy, sick* is a paradigmatic example: one can be healthy or sick with respect to various dimensions, such as blood pressure, cholesterol or blood sugar level. By contrast, there is but one dimension associated with an adjective like *tall* (i.e. height).

An available test for multidimensionality is the admissibility of "dimension-accessing" operators and modifiers, such as the PPs *with respect to ...* and *in some/most/every respect(s)*. Compare *healthy* (multidimensional) to *tall* (dimensional):

- 24) Natalia is healthy in some/most/every respect(s) / ... with respect to blood pressure, but not cholesterol.  
 25) # Matheus is tall in some/most/every respect(s) / ...with respect to ?

Moral adjectives are multidimensional, as shown by the fact that they also admit such "dimension-accessing" operators and modifiers:

- 26) What she did was despicable with respect to its level of cold-bloodedness.  
 27) Natalia is virtuous in some/most/every respect(s).

Importantly, multidimensional adjectives admit maximum/minimum standard modifiers like *perfectly/slightly*. But when they combine with multidimensional adjectives, these modifiers do not reference endpoints on a scale, but rather, they quantify over the dimensions associated with the relevant adjectives: *perfectly healthy* means, roughly, *healthy in all respects/dimensions*; and *slightly sick* means *sick in some respect/dimension* (see Sassoon 2016, p. 10). Furthermore, those modifiers are not interchangeable: *perfectly* only admits multidimensional adjectives that have a *universal* interpretation, that is,

adjectives for which the bare positive form requires that the individual is above the relevant threshold for *every* dimension. By contrast, a modifier like *slightly* is only acceptable with multidimensional adjectives that are interpreted *existentially*, that is, where the bare positive form only requires for the individual to be above the threshold for *some* dimension. *Healthy* and *sick* are universal and existential, respectively (Sassoon 2016, p. 3, proposes the generalization that positive multidimensional adjectives tend towards universal interpretations, while negative adjectives tend towards existential interpretations). Thus, the following pattern emerges:

28) perfectly / ?? slightly healthy

29) ?? perfectly / slightly sick

Thus, we submit that that is how the acceptable APs in (22)–(23) are to be interpreted: *perfectly ethical* means *ethical in every respect*; *slightly unethical* means *unethical in some respect*. Importantly, these interpretations are consistent with the claim that such adjectives are relative.

This suggests the following modification of the lexical entries that we started out with. The semantic value of a multidimensional and *universal* gradable expression  $G+$ , in its positive form, and according to the multidimensionality hypothesis, would now be

$$\llbracket \alpha \text{ is } G+ \rrbracket^{M,w,d} = 1 \text{ iff } \forall d \in \Delta_G, \mu_G^d(\llbracket \alpha \rrbracket^{M,w}) \geq \delta_G^d,$$

while the semantic value of a multidimensional and *existential* gradable expression  $G-$  (in its positive form) would be

$$\llbracket \alpha \text{ is } G- \rrbracket^{M,w,d} = 1 \text{ iff } \exists d \in \Delta_G, \mu_G^d(\llbracket \alpha \rrbracket^{M,w}) \geq \delta_G^d,$$

where  $d$  is a parameter ranging on the set of relevant dimensions  $\Delta_G$  (we refer the reader to Sassoon 2016 for more details).

Thus, we hypothesize that moral adjectives are multidimensional, relative gradable adjectives whose threshold is not sensitive to modification by comparison class.

Before we move on however, one may wonder: why venture such generalization? Why not entertain the possibility that, for instance, some moral adjectives are absolute-standard? There is an empirical and a theoretical reason for this: on empirical grounds, we simply have yet to come across a moral adjective that shows a markedly

different semantic behavior from what we observe here. Our choice of examples in this paper is not selective; we have picked moral adjectives at random and observe that they reveal a – relatively stable – semantic pattern. On theoretical grounds, we consider it a plausible hypothesis that all moral adjectives (and possibly other evaluative adjectives as well) have similar semantic properties due to their sharing certain mathematical or structural properties inherent to values in general. For instance, it's plausible that orderings of objects according to their moral value are only partial. Thus, one might reasonably expect that all moral adjectives give rise to partial orderings as well (thereby allowing for incomparabilities). However, we think that if the foregoing hypothesis were true, that would constrain the mathematical properties of the scales of these adjectives (that is, the properties of the measure function  $\mu$ , see next section), but not necessarily the properties of their thresholds. In other words, that moral adjectives track mathematical properties of values is compatible with moral adjectives being relative or absolute standard. So we take there to be theoretical reasons to expect some uniformity at the level of the internal scale of these adjectives, and empirical – but of course defeasible – reasons to expect uniformity at the level of thresholds for the positive form.

### 3. A FORMAL PERSPECTIVE ON SCALES

In this section we offer an additional argument for our thesis, partially independent from the preceding section, as we slightly change our point of view. Instead of taking a “bottom-up” approach, as it were, starting from linguistic tests, we now adopt a “top-down” perspective. Even if the fact that a moral adjective is multidimensional allows for each dimension to have different kinds of scale, plausibly the dimensions have to be aggregated or combined into a single scale, and we can still explore what are the properties of *that* scale based on the observable linguistic properties of the adjective itself.

This is what we now turn to, namely a more mathematically oriented discussion of the properties of scales for moral adjectives, i.e. we venture into the properties of  $\mu$ . We show results of two classes: first, boundedness properties are compatible both with ratio and interval scale and are, therefore, uninformative; second, we prove that

moral adjectives cannot always have a ratio scale because of some non-additive cases. These results, we submit, are a further argument for the multidimensionality hypothesis.

For concreteness, we consider a structure of individuals, an ordering, and a sum operation:  $S = (S, \supseteq, \sqcup)$ . We seek to define an order-preserving mapping  $\mu$  from  $S = (S, \supseteq, \sqcup)$  into  $(\mathbb{R}, \geq, *)$ , where  $*$  might be addition or a more complicated operation. Such a mapping is called a *representation* of  $S$ ; to show there is such a function is to prove a representation theorem.

There are three kinds of scales used in empirical and social sciences (barring *absolute* or *nominal* “scales”, which just label elements without any kind of quantitative ordering or measurement – see Duncan and Narens (1987) for a concise but high-level survey):<sup>7</sup> *Ordinal scales* represent the ordering among the elements to be measured with the usual ordering among (real) numbers, with no further properties assumed: we do not know anything about the respective distances between elements. Examples of ordinal scales are the scales that we come across in surveys, such as those ranging from “1 = very boring” to “5 = very exciting”. More precisely,  $\mu$  is an *ordinal scale* iff if for all  $a, b \in S$  and  $a \supseteq b$ , then  $\mu(a) \geq \mu(b)$ . *Interval scales* represent the ordering with the usual ordering among (real) numbers, but where differences are meaningful, i.e. differences represent actual distances between the elements to be measured. Examples of such scales are the Celsius or Fahrenheit scales. More formally,  $\mu$  is an *interval scale* if the following conditions are met: (i) if  $a, b \in S$  and  $a \supseteq b$ , then  $\mu(a) \geq \mu(b)$ ; (ii)  $\mu(a \sqcup b) = k\mu(a) + p\mu(b) + q$ , with  $k, p, q \in \mathbb{R}^+$ ; (iii) for any  $\mu'$  satisfying (i) and (ii), there are  $n, m, k', p', q'$  with  $n, k', p', q' \in \mathbb{R}^+$ ,  $m \in \mathbb{R}$  s.t.  $\mu'(x) = n\mu(x) + m$ ,  $q' = nq + m(1 - k - p)$ ,  $k' = k$ ,  $p' = p$ , i.e. an interval scale is unique only up to positive affine transformation. The following known theorem lists the conditions on  $S$  for the existence of an interval scale (see Krantz et al. 1971, p. 294ss.):

**Theorem 1** Let  $S = (S, \supseteq, \sqcup)$  be a structure such that for all  $a, b, c, d, e, f$ , it is monotonic ( $a \supseteq b$  iff  $a \sqcup c \supseteq b \sqcup c$ ), bisymmetric ( $(a \sqcup c) \sqcup (c \sqcup d) \approx (a \sqcup c) \sqcup (b \sqcup d)$ ), restrictedly solvable, Archimedean and  $\supseteq$  is a weak ordering. Then  $\mu$  is an *interval scale*.

---

<sup>7</sup> Although an infinite number of different scales can be characterized more precisely with reference to homogeneous and point-uniqueness features.

Finally, *ratio scales* represent the ordering with the usual ordering among real numbers, where difference and multiplication are meaningful so that ratios are preserved. Examples of such scales are the weight or height scales. More precisely,  $\mu$  is a *ratio scale* if the following conditions are met: (i) if  $a, b \in S$  and  $a \supseteq b$ , then  $\mu(a) \geq \mu(b)$ ; (ii)  $\mu(a \sqcup b) = \mu(a) + \mu(b)$ ; (iii) for any  $\mu'$  satisfying (i) and (ii), there's an  $n \in \mathbb{R}^+$  s.t.  $\mu'(x) = n\mu(x)$ , i.e. a ratio scale is unique only up to linear transformation. The following known theorem lists the conditions on  $S$  for the existence of a ratio scale (see Krantz et al 1971):

**Theorem 2** Let  $S = (S, \supseteq, \sqcup)$  be a structure such that it is *positive*, *monotonic*, *solvable*, *Archimedean* and  $\supseteq$  is a weak ordering. Then  $\mu$  is a *ratio scale*.

We moreover say that if a scale obeys condition (ii) in the previous paragraph for all distinct  $a, b \in S$ , it is *additive*. Ratio scales are therefore additive. Interval scales are not additive, in that the “sum” of two values lies in the middle. Weighted averages are an example of such an operation.

One naturally wonders whether the scale of moral adjectives is among these options. At the very outset, we can exclude the idea that moral adjectives have an ordinal scale if we assume they are at least minimally more structured than a mere ordering. Lassiter (2016) offers an argument in this sense. Thus, in the following we consider three data points in favor of either ratio or interval scales: boundedness, modification with ‘twice’, and inferences about concatenation.

First, boundedness. In the previous section we considered the possibility that pairs of moral adjectives were partially closed. Can boundedness properties reveal something about scale structure? Unfortunately, no. The relationship of interval and ratio scales with boundedness is summarized in the following two theorems:

**Theorem 3** Let  $\mu$  be a *ratio scale*. Then it can be open; lower-bounded; upper-bounded; fully closed.

**Proof** (Sketch) We construct an example of each. Take a measure which respects the usual probability axioms. We check that it respects conditions (i)–(iii) of ratio scales, therefore it is a ratio scale. Normally, it is fully closed, i.e. its range is  $[0,1]$ . Remove one or both of the

endpoints and one gets the following cases:  $(0,1],[0,1),(0,1)$ . We check that each measure, so modified, respects conditions (i)–(iii).

**Theorem 4** Let  $\mu$  be an *interval scale*. Then it can be open; lower-bounded; upper-bounded; fully closed.

**Proof** By analogy to the above.

Therefore, we can't exclude the possibility that moral adjectives have either a ratio or an interval scale depending on the hypothesis that they are open or partially or fully closed. Nonetheless, a second data point may be the admissibility of modifiers such as 'twice': when employed in certain contexts, 'twice' appears to point toward a ratio scale, since there is explicit talk of additions or multiplications (see Lassiter 2016). And indeed, 'twice' is acceptable as a modifier of moral adjectives; see the following examples in combination with 'cruel':

- 30) You're unwilling to buy a mousetrap, but you're happy to buy a sticky platter? [...] If anything, it's twice as cruel as a mousetrap. At least the mouse is killed instantly. Imagine the slow agonising death of being stuck to a plate.<sup>8</sup>
- 31) Well you've taken the best of our sailors, and, You've taken my love from me. [...] The sea is twice as cruel.<sup>9</sup>
- 32) If you are guilty or something, just break up. Telling is twice as cruel.<sup>10</sup>
- 33) If anger is cruel, then jealousy is (being) twice as cruel.<sup>11</sup>

However, the use of *twice*, by itself, is not conclusive evidence. For all these sentences show, speakers could be speaking loosely or metaphorically. Moreover, the lack of any standard measure of moral value impedes any precisification of what anyone could mean by describing something as 'twice as cruel' as something else.

<sup>8</sup> <http://www.bangkokpost.com/print/1037653/>

<sup>9</sup> Song *Oh Cruel Sea!*, by Jonny Nutt.

<sup>10</sup> <http://whisper.sh/whisper/05338d292f99f2736bfbf12c4361ad8c88ecd1/If-you-are-guilty-or-something-just-break-up-Telling-is-twice-as-cruel>

<sup>11</sup> Ashish Raichur, *Laying the Axe to the Root*, p. 33. Quite interestingly, this example opens up the possibility of investigating higher-order scales, i.e. scales introducing comparisons of properties (according to a higher-order property), rather than of individuals. For this reason we employ this example with some hesitation.

Finally, we should also take into account facts directly about concatenation. In particular, the following fact holds:

**Theorem 5** If  $\mu$  is not additive, then it is not a ratio scale.

It is enough to show that moral adjectives are not (always) additive to show that they do not have a ratio scale.

Consider the following example:

34) It is cruel to make him run in the heat.

35) It is cruel to make him run in the rain.

36) \*It is twice as cruel to make him run in the heat and in the rain.

(36) is plausibly unacceptable as an inference from (34) and (35), for the concatenation of the two factors (heat and rain) might in fact make the run pleasant, or at any rate not nearly as cruel as if either factor was present singly, (provided we understand (36) as a fusion of the two situations, rather than a mere serial repetition),<sup>12</sup> so that even the following becomes unacceptable as a consequence of (34) and (35):

37)\* It is cruel to make him run in the heat and in the rain.

The concatenation of the two factors, in this case, would even be less cruel, if at all, than either factor taken alone.

The choice of ‘cruel’ was somewhat arbitrary: a similar reasoning seems to apply to thin adjectives as well. We thus suggest that moral adjectives do not have ratio scales. However, we cannot conclude yet that moral adjectives have an interval scale, since in fact, further conditions need to be met for them to do so. Lassiter 2016 argues that ‘good’ has such a scale essentially for abductive reasons. But both the linguistic data and theorems 3, 4, and 5 are not enough to conclude that all moral adjectives have an interval scale. In fact, there are infinitely many different scales to choose from. At present, we have to leave open two other possibilities as well: *first*, moral adjectives do not

---

<sup>12</sup> A possible objection may identify the reason for the fact that (36) does not follow from (34) and (35) not in the scalar properties (or absence thereof) of cruelty but in the non-standard behavior of the conjunction in this particular case: “to make him run in the heat and in the rain” would not be equivalent to “to make him run in the heat and to make him run in the rain.” While we don’t think this is the case, such concerns are immaterial to our point to the extent that there is a mechanism to talk about the “fusion” of the two situations. We thank an anonymous referee for pressing on this point.

have an interval scale, but rather yet another scale, perhaps with a very complicated structure, or at any rate, a non-standard structure; *second*, moral adjectives may not have a unique scale at all (even up to the appropriate notion).

Both these possibilities would be explained by the hypothesis that moral adjectives are multidimensional, if their multidimensionality is constructed in a way compatible with associating possibly different scales to each dimension.

#### 4. CONCLUSION

In this paper, we've (preliminarily) observed that moral adjectives are multidimensional, relative-standard adjectives. More tests, and especially, more experimental results should be obtained, in order to decide between the alternative features of scales presented here. We leave this for future work, but we note that getting clear on the scale of moral adjectives has important consequences for ethical theory, provided that moral language is somewhat indicative to ethics. Of course, one may argue that linguistic evidence is a poor indicator for philosophical analysis, and the structure of these concepts is ultimately a matter of normative philosophical theory, as Erich Rast (p.c.) notes, and that's why linguistic data may seem generally inconsistent. A natural way of reconciling the apparent inconsistency of the examples, however, is to allow for many different measures which are to be specified on a case-by-case basis. Thus, there would be different *kinds* of scales (as opposed to different scales equivalent up to some notion to be specified): sometimes the scale would be additive, and sometimes not, sometimes it would be ratio-like, sometimes interval-like. This scenario fits well with moral particularism. The links between multidimensional approaches and particularist approaches to ethical theory are, to the best of our knowledge, still unexplored.<sup>13</sup>

---

<sup>13</sup> The authors would like to thank the audience at the PhilLang 2017 conference, Natalia Karczewska, Erich Rast, Isidora Stojanovic, Matheus Valente and two anonymous referees. Andrés Soria Ruiz was supported by Obra Social La Caixa and grant numbers ANR-10-LABX-0087 IEC and ANR-10-IDEX-0001-02 PSL.



## REFERENCES

- Barker, C. (2002). The Dynamics of Vagueness. *Linguistics and Philosophy*, 25(1), 1–36.
- Bierwisch, M. (1989). The Semantics of Gradation. In M. Bierwisch and E. Lang (Eds.), *Dimensional Adjectives* (pp. 71–261). Berlin: Springer-Verlag
- Chang, R. (1997). Introduction. In R. Chang (Ed.), *Incommensurability, Incomparability and Practical Reason*. Cambridge, MA: Harvard University Press.
- Cresswell, M. (1976). The Semantics of Degree. In B. Partee (Ed.), *Montague Grammar* (pp. 261–292). New York: Academic Press.
- Darwall, S., Gibbard, A., Railton, P. (1992). Toward Fin de Siècle Ethics: Some Trends. *The Philosophical Review*, 101(1), 115–189.
- Fara, D. G. (2003). Gap Principles, Penumbral Consequence, and Infinitely Higher-Order Vagueness. In J. C. Beall (Ed.), *New Essays on the Semantics of Paradox*. Oxford University Press.
- Kennedy, C. (2007). Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives. *Linguistics and Philosophy*, 30(1), 1–45.
- Klein, E. (1980). A Semantics for Positive and Comparative Adjectives. *Linguistics and Philosophy*, 4(1), 1–45.
- Kölbel, M. (2003). Faultless Disagreement. *Proceedings of the Aristotelian Society*, 104(1), 53–73.
- Krantz, D. et al. (1971). *Foundations of Measurement*, Vol. 1. New York: Academic Press.
- Kratzer, A. (2012). *Modals and Conditionals: New and revised perspectives*. Oxford: Oxford University Press.
- Laserson, P. (2005). Context Dependence, Disagreement, and Predicates of Personal Taste. *Linguistics and Philosophy*, 28(6), 643–686.
- Lassiter, D. (2016). *Graded Modality. Qualitative and Quantitative Perspectives*. Oxford University Press.
- Lassiter, D., Goodman, N. D. (2013). Context, Scale Structure, and Statistics in the Interpretation of Positive-Form Adjectives. In Todd Snider (ed.), *Semantics & Linguistic Theory (SALT) 23* (pp. 587–610). Ithaca: CLC Publications.
- Lassiter, D., and Goodman, N. D. (2015). Adjectival Vagueness in a Bayesian Model of Interpretation. *Synthese*, 1–36.
- Liao, S.-Y., Meskin, A. (2015). Aesthetic Adjectives: Experimental Semantics and Context-Sensitivity. *Philosophy and Phenomenological Research*, 94(2), 371–398.
- Liao, S.-Y., McNally, L., Meskin, A. (2016). Aesthetic Adjectives Lack Uniform Behavior. *Inquiry*, 59(6), 618–631.
- Luce, R. D., Narens, L. (1987). Measurement Scales on the Continuum. *Science*, 236: 1527–1532.
- MacFarlane, J. (2014). *Assessment Sensitivity: Relative Truth and Its Applications*. Oxford: Oxford University Press.
- Meier, C., van Wijnbergen-Huitink, J. (Eds.) (2016). *Subjective Meaning: Alternatives to Relativism*. Berlin: De Gruyter.

- Sassoon, G. W. (2016). Multidimensionality in the Grammar of Gradability, MS.
- Stephenson, T. (2007). Judge Dependence, Epistemic Modals, and Predicates of Personal Taste. *Linguistics and Philosophy*, 30(4), 487–525.
- Stojanovic, I. (2007). Talking About Taste: Disagreement, Implicit Arguments, and Relative Truth. *Linguistics and Philosophy*, 30(6), 691–706.
- Stojanovic, I. (2016). Expressing Aesthetic Judgments in Context. *Inquiry*, 1–23.
- Temkin, L. (2012). *Re-thinking the Good*. Oxford: Oxford University Press.
- Umbach, C. (2016). Evaluative Propositions and Subjective Judgments. In C. Meier and J. van Wijnbergen-Huitink (Eds.), *Subjective Meaning: Alternatives to Relativism* (pp. 127–168). Berlin: De Gruyter.