DOI: 10.1111/nous.12454

ARTICLE

Noûs

Dimensions of Value

Brian Hedden[1] | **Daniel Muñoz**[2]

[1]Australian National University, Australia

[2]University of North Carolina at Chapel Hill, Chapel Hill, NC, USA

Correspondence
Daniel Muñoz, Australian National University, Australia.
Email: Daniel.munoz@unc.edu

Abstract

Value pluralists believe in multiple dimensions of value. What does betterness along a dimension have to do with being better overall? Any systematic answer begins with the Strong Pareto principle: one thing is overall better than another if it is better along one dimension and at least as good along all others. We defend Strong Pareto from recent counterexamples and use our discussion to develop a novel view of dimensions of value, one which puts Strong Pareto on firmer footing. We conclude by defending Dimensionalism, the hypothesis that overall value relations are determined solely by how things compare along value dimensions. These are first steps towards a more systematic value pluralism.

1 | INTRODUCTION

Value monists, such as the hedonist who values only pleasure, believe there is exactly one way in which things can be good. *Value pluralists* believe in multiple ways of being good, or, in other words, multiple dimensions of value.

The value pluralist faces two fundamental questions. First, what is a dimension of value? Here we think pluralists should say more; they often do not explain what they have in mind, and sometimes use 'dimension' in conflicting ways. Second, how, if at all, do dimensions of value combine to determine which things are better than which *overall*? Here we think pluralists *must* say more. Traditionally, pluralists have not sought a fully systematic theory—a finite set of exceptionless principles linking dimensions to overall values.[1] But in recent work, even the most minimal

[1] See, e.g., Nagel (1979: 131) and Wolf (1992: 785).

The authors are listed alphabetically and contributed equally to this article.

linking principles have come under attack (Chang, 2002a, 2016a; Constantinescu, 2012; Hsieh, 2005), endangering the prospects for even a minimally systematic pluralism.

In this paper, we take up the cause of the systematic pluralist. We begin with the question of how values combine, defending the principle known as *Strong Pareto* against recent counterexamples (§2-5). We then draw from this discussion to develop a conception of dimensions of value as distinct respects in which a thing can be good or bad, to which overall value is responsive (§6). We argue that our conception, in turn, yields a better argument for Strong Pareto than the more common "No Objection Argument," which justifies an even stronger version of Strong Pareto that gives rise to cyclic betterness (§7). We close with a fresh question: does overall value depend solely on how things stand on the relevant dimensions of value? We think an affirmative answer should be a working hypothesis for any value pluralist in search of a more principled theory (§8).[2]

2 | THE PARETO PRINCIPLES

On any systematic value pluralism, a thing's overall value should be responsive to its value along the relevant dimensions of value.[3] Responsive in what way?

A minimal answer is that improving things along dimensions should improve things overall, other things being equal. From this we get two principles, the first of which entails the second:[4]

Strong Pareto
If A is better than B along some dimension and at least as good along all others, then A is overall better than B.

Weak Pareto
If A is better than B along every dimension, then A is overall better than B.

The Pareto principles may seem obvious, even trivial. We think they are the bare minimum for any principled pluralist, but not everyone agrees. Some say that merely improving a thing along some dimensions, holding everything else fixed, can leave it overall no better—or even worse. This is a striking claim. If it is correct, then Strong Pareto is false. Even worse, their arguments to this effect would put even Weak Pareto in jeopardy, leaving the prospects for systematic pluralism very grim indeed. Dimensions of value would combine in ways that seem not just complex but unintelligible. In the next sections (§§3–5), we defend the Pareto principles against three kinds of counterexamples, involving organic unities, golden means, and tiny improvements.

[2] Some pluralists, like Dancy (2004) and Ross (1930: 41), are skeptical of *any* principles for resolving moral conflicts. But our principles are so minimal that they do not apply in conflicts; even Ross and Dancy would be loath to give them up.

[3] Compare this with Hurley's (1989: 228) claim that "criteria are directional, such that any one alternative's ranking above a second with respect to a criterion counts per se in favour of the first." (In her terminology, 'criterion' means 'value dimension'.)

[4] We adapt our Pareto principles from social choice theory, where they are formulated more narrowly as principles that link overall moral betterness to what is prudentially better for individuals. (In other words, individual welfare levels play the role of dimensions.) For an introduction to social choice, see Gaertner (2009) or Sen (2017).

3 | ORGANIC UNITIES

Ruth Chang argues that an option can be made worse overall by improving it along one dimension, leaving it just as good along every other, when this creates a kind of imbalance between dimensions. She illustrates with an example:

> …one life, identical to another in all respects except better with respect to pleasure, may not be the better life because too much pleasure may make for an unbalanced life; what makes a life best may be an organic unity of various goods. (2002b: 674)

According to Chang, Strong Pareto fails here because of "organic unities." Improving something in one respect need not improve it overall, because how good it is overall depends on the total balance of the respects. Improving a respect might upset the balance. "A small Pareto improvement," as Chang (2002a: 132) puts it, "does not always make an item at least as good as its unimproved counterpart: it might make it worse." (A "Pareto improvement" leaves a thing better along one dimension and at least as good along all others.)

Chang also has other examples, three of which are relevant here:

I. adding a murder-mystery epilogue to *Pride and Prejudice*—resulting in more literary merit in a respect, but a worse novel overall (2016a : 402);
II. doubling the most pleasant (intense!) four seconds of a 90-second roller coaster ride, with the result being a ride that is less pleasurable overall (2016a : 401–2);
III. making a philosopher better with respect to "historical sensitivity" might not result in a better philosopher, but rather a historian, whose talents may not even be comparable to the original philosopher's in terms of overall betterness, worseness, and equality. (2002a : 22)

All of these cases are purported counterexamples to Strong Pareto that appeal to the importance of balance or unity. A Pareto improvement upsets the balance, thereby making things worse.

Note that Chang's appeal to balance would also seem to undercut Weak Pareto. For just as we could upset balance by improving something along one dimension and leaving it unchanged along all others, we could also upset balance by improving something greatly along one dimension and improving it only slightly along all others. Insofar as the importance of balance can make the former a change for the worse, it should likewise be able to make the latter a change for the worse.

Let's grant that Chang is on to something: balance matters. But that is not a good reason to reject Strong or Weak Pareto. To see why, we need lay out her example of the overly pleasurable life a bit more precisely. Suppose Chang's actual life, L1, is good in the respects that matter, and, for simplicity, imagine that there are only two such respects: pleasure and meaning. We might assign L1 a value score of (50, 50), where the two numbers measure pleasurableness and meaningfulness, respectively. Now compare this to L2, which gets a score of (100, 50), indicating that it has more pleasure and the same level of meaning. On Chang's view, L2 might be worse than L1 in virtue of L2's imbalance—with a lopsided 100 in pleasure—even though L2 is no worse in either dimension. It is this extreme conclusion that conflicts with Strong Pareto.

In particular, Strong Pareto conflicts with the conjunction of two claims: first, that L2 is a Pareto improvement over L1; and, second, that L2 is nonetheless overall worse because its imbalance. But these claims, unlike Strong Pareto, are easy to deny.

Start with the first claim: that L2 is a Pareto improvement. On Chang's own view, L2's imbalance makes it overall worse than L1. So why not say that L2 is indeed worse along a dimension—namely, balance? We already have dimensions for pleasure and meaning. We could simply add a third dimension to measure balance. Suppose, for example, that balance is equal to –2 times the absolute difference between the other two dimensions. From this, we get new overall values for L1 (50, 50, 0) and L2 (100, 50, –100). On such an analysis, L2 could still be worse than L1 overall, but, since L2 is no longer a Pareto improvement due to its bad balance, Strong Pareto is safe. This take on the case perfectly fits the Pareto principles, and Chang has no argument against it.[5]

But suppose she finds such an argument, and we have to give up the idea of balance as a dimension of value.[6] Then we can and should deny the second claim: that the imbalanced L2 is overall worse than L1. For why *should* L2 be worse, if it is far better in one way, and no worse in any way? If imbalance does not worsen a life in any dimension whatsoever, it is unclear how it could detract from the overall value of a life.

To be sure, balance might still *matter* for the value of a life, even if it is not itself a dimension of value. But the familiar views on which this this is true do not conflict with the Pareto principles. For example, on some views, dimensions each count the same but have diminishing marginal contributions to overall value.[7] Balance matters in the sense that improvements along the "worst-off" dimensions yield more overall value than equal-sized improvements along the "best-off" dimensions: if your life is a meaningless joy, a boost to meaning will count for more than a similar boost to pleasure.[8] But this view still says that, other things being equal, it's always better overall to boost a dimension, even if doing so leads to an imbalanced life. So the view still satisfies Strong Pareto. We conclude that balance and organic unities do not threaten the Pareto principles.

4 | GOLDEN MEANS

Chang also argues that Pareto improvements can make something overall worse if it already has "too much" value in a dimension. She writes:

[5] This proposal is analogous to *Egalitarianism*, the view that it is valuable in itself if people are equally well off (Parfit, 1997: 204; Temkin, 1983). Balance among dimensions is like equality between different people's levels of welfare.

[6] To argue that balance is not a dimension, one might appeal to a *Leveling Down Objection* (analogous to an objection used against Egalitarians; see Parfit (1997: 211) and Temkin (1983)). Suppose that balance is a dimension, and consider an imbalanced life L2 (100, 50, –100). There would then be two ways to improve balance: "level up" by making the life better with respect to meaning, or "level down" by making the life worse with respect to pleasure. But this is absurd, the objection goes: "leveling down" cannot make a life better in any respect. Whatever the merits of this objection, Chang would not appeal to it. On her view, leveling down L2 to get L1 improves the life *overall*. This is even more extreme than saying that leveling down improves L2 with respect to balance.

[7] In particular, we have in mind a view like the following. Suppose X is overall better than Y iff the sum of the square roots of X's scores along each dimension is greater than the sum of the square roots of Y's scores along each dimension. On this view, L3 (49, 49) would be better overall than L4 (100, 0) since 14 is greater than 10, even though L3's dimensions add up to only 98 as opposed to 100. (This view assumes *inter alia* that dimensions are ratio-scale measurable, that we can make interdimensional comparisons, and that objects receive only non-negative scores along any given dimension.) Compare Hurka (1993, ch. 7) on the well-rounded life.

[8] This proposal is analogous to *Prioritarianism*, the view that boosts to a person's welfare matter more when that person is, in absolute terms, worse off (Parfit, 1997: 213; Temkin, 1983). Priority for the "worst-off" dimensions is like priority for the worst-off people.

> [M]ore is not always better; you can have too much of a good thing. A manifestation of too much of a particular quality of raw philosophical talent may be worse with respect to philosophical talent than a lesser amount of that raw quality. Some values are structured so that, beyond a certain threshold, more of a value as manifested in an aspect is worse. In general, we should not expect the trade-offs among value aspects to behave in a tidy, monotonic fashion. (2002a : 11)

Chang says one can have "too much" talent. As we interpret her, she means that being *too good* along the dimension of talent makes a philosopher *worse overall*, even if other dimensions are held equal.[9] This is like the case of organic unities, except that we have a goodness glut in one dimension, not an imbalance between several. We could have such a glut even if there were just a single relevant dimension of value—even if value monism were true—whereas balance only makes sense if there are multiple dimensions of value.

Is Chang right about thresholds? Can there be "too much" goodness in a dimension? In our view, the example of talent is not intuitively plausible, especially when we consider that similar reasoning can lead to failures of Weak Pareto. Suppose an option is at the threshold along each dimension of value. Then we can improve that option along each dimension to make it overall worse. Consider a philosopher with "too much" talent, originality, clarity, and so on. For Chang, this philosopher is overall worse than someone with a moderate amount of goodness on each dimension, despite the "worse" philosopher being vastly better in every relevant way—a paradoxical result.

There is also, as before, a plausible way to reconcile the Pareto principles with Chang's intuitions by rethinking what the relevant dimensions are. Suppose Chang is right that while it is bad to have no talent, one can also have "too much." Then there is a "golden mean" for talent, and the closer one is to the mean, the better one is overall. Shouldn't we then say that the relevant dimension of value is not *amount of talent*, but rather *proximity to the ideal amount?*

Consider an analogy from virtue ethics. If courage is the measure of one's willingness to take risks, then plausibly there is a threshold—a golden mean—below which more courage improves overall virtue, and above which more courage worsens overall virtue. (Recklessness is a vice, after all.) But no virtue ethicist would conclude that one can sully one's virtue simply by overly "improving" in courage. Instead, the right thing to say is that amount of "courage," so defined, is not a dimension of virtue. Rather, the relevant dimension is *proximity to the ideal amount of courage*, which is lacking in both the reckless and the cowardly.[10]

The takeaway is that true dimensions do not have "golden means," but "proximity to the golden mean" can be a dimension. This point holds for dimensions of virtue, value, and any other multidimensional property. If being a good latte consists partly in having a certain ratio of milk to coffee, then ratio of milk to coffee is not itself a dimension of goodness qua latte; the real dimension is instead closeness to the ideal ratio.

[9] Chang's claim, in her terms, is that things can be worse with respect to a covering value if they manifest too great a quantity of value in a contributory aspect (2002a : 11–4; see also Hsieh (2005: 186n15)). Elsewhere, she makes it clear that "contributory aspects" are indeed "dimensions," and that her target is Strong Pareto: "sometimes one thing can be better than another along a single dimension, without differing from it in [*sic*] along any other dimension, but still be not better overall" (2016a : 408).

[10] Aristotle (1999), by contrast, would define courage itself as one's proximity to the ideal level of willingness to take risks. (Here we simplify *Nicomachean Ethics* III.6-9).

If we see value dimensions in this way, we will not be moved by examples such as Chang's, where "talent" or "courage" is supposed to be a dimension with a golden mean. If it has a golden mean, it is not a dimension—but "proximity to the golden mean" might be. So golden means are no problem for the Pareto principles.

5 | TINY IMPROVEMENTS

Chang, we have seen, thinks that Pareto improvements can lead to "too much goodness" in two distinct ways. In the case of organic unities, having too much value in one dimension relative to the others can create a bad imbalance; in the case of golden means, too much value in a dimension absolutely makes a thing overall worse. But in each case, we can avoid the problem by rethinking the relevant dimensions. If we imbalance makes a thing worse, we should count balance as a dimension; if there is an ideal amount of talent, we should think of the relevant dimension as proximity to that ideal.

Next comes a new problem for Pareto. Instead of "too much goodness," the worry is "not enough betterness."

Constantinescu suggests that an increase along one dimension of value might be *too small* to amount to an overall increase in value, even holding fixed every other dimension. He writes that:

> …our comparative predicates almost always allow for certain margins for error. To employ the popular metaphor of weighing, we might say that the scales we use to weigh goods and rank them against one another are not sufficiently sensitive to register very slight differences or changes in value. Imagine, for instance, that someone were to add two mildly insightful lines to one of David Lodge's novels. Would the resulting book be better than the original? The issue appears indeterminate: it would seem that there are always certain amounts of value that are too small to tip our scales. (2012 : 59)

Now, Constantinescu isn't just saying that small changes to a book can fail to add any value. The book *would* be more insightful with the two extra lines.

Constantinescu's point is rather that small "changes in value" along a dimension can leave a thing's overall value equal to what it was before—so, a tiny boost to one dimension may not make a thing better overall, even with other dimensions left the same. Adding lines to the novel fails to make it better *despite* making it a little more insightful and otherwise equally good. This violates Strong Pareto.[11]

Constantinescu's example, like Chang's examples, may also threaten Weak Pareto. For suppose we improve the novel along *each* dimension of value, but just barely, so that its overall value is unchanged. (Perhaps, for each dimension, we tweak a line of the novel to make it better along that dimension.)[12]

[11] This is how we read Constantinescu, though admittedly he is not explicit in concluding that we should reject the Pareto principles, for his interests lie elsewhere. As we explain in the following footnote, Hsieh (2005) uses similar considerations to explicitly motivate a rejection of the Pareto principles.

[12] Hsieh (2005) would reject Weak Pareto for a similar reason. He thinks that, in some domains, our value judgments are "clumpy" ("low resolution"), ignoring tiny differences along dimensions. He illustrates with the example of "clumping" papers by grade:

We think Constantinescu is on to something, but again we see no need to reject Strong or Weak Pareto. The problem is that he treats tiny changes as being, in effect, no change at all. Why do that? A tiny improvement might not be notable, or even noticeable, but it should still count as an improvement.

Constantinescu might reply that *true* improvements must "tip our scales"—that is, they must be big enough for us to "register." But if the tiny tweak to the novel is too small to register, how could it improve the novel with respect to insight? Doesn't the boost to insight show that the change can "tip our scales?"

The counterexample thus seems unstable. If a tiny change can be a Pareto improvement, however tiny, there is no clear reason why it cannot also be an overall improvement, however tiny.

The moral, this time, is not about identifying the right dimensions. Instead, the moral is that even puny Pareto improvements make things better overall.

6 | DEFINING DIMENSIONS

Let us now turn from tiny differences to the bigger picture.

We have defended the Pareto principles against three supposed counterexamples: organic unities, golden means, and tiny improvements. In each case, the argument against the Pareto principles seemed to rest on a questionable view about what the relevant dimensions are, or how dimensions relate to overall value.

Now let us turn from intuitions about examples to the deeper theoretical issue: *what are dimensions of value?* Our discussion so far suggests an answer: a dimension is (i) a genuine and distinct respect in which things can be good or bad,[13] (ii) to which overall value is responsive.[14] Both parts of the definition matter, so let us spell them out.

The first part—that dimensions are genuine and distinct respects in which things can be good or bad—lets us say that things *correlated* with dimensions might not themselves be dimensions. We can illustrate this with the example of golden means. The amount of courage one has correlates with how close one is to the ideal amount. But, in our view, it is only *proximity to the ideal amount* that is a dimension of value; raw courage, as distinct from this, is not itself a genuine dimension.[15]

> although we might be able to see that Ahmed's "B" article is better than Anne's "B" article in all of the relevant respects that comprise the covering consideration, given the resolution of the comparison it simply is not meaningful to place the two articles in different clumps. They are equally 'B' articles, and therefore, equally good (2005 : 186).

This violates Weak Pareto, a fact that Hsieh understates. (As he puts it, clumping leads to "dominance insensitivity," which occurs when A is at least as good as B on all dimensions and no better overall. But such "insensitivity" is nothing special; notice that A is at least as good as itself on all dimensions, and yet A is not overall better than itself.)

[13] Elsewhere, one of us calls (i) "the dimensionality constraint" and endorses it without argument: "there should be exactly one dimension of value for each genuine and distinct way in which things can be good or bad" (Muñoz, forthcoming: 18).

[14] For a related idea, consider Nebel's (2015: 11) definition: value dimensions are "respects or ways in which [objects] might be valuable," which have at least "ordinal structure." (That is, if F is a value dimension, then one thing can be F-er than another.) We agree with Nebel about ordinal structure, though we want to say more about "respects" and how they relate to overall value.

[15] When can we be sure that we are dealing with two dimensions? A telltale sign: one can have something to *regret* about choosing an option even if that option is not overall worse than the alternative. (See, e.g., De Sousa, 1974: 536; Chappell, 2015: 328.)

The second part of the definition—that overall value is responsive to dimensions—lets us say that improving something along a dimension will always, other things being equal, improve it overall (see §6). We can illustrate this with the example of tiny improvements. If tweaking two sentences of a novel makes it better with respect to insight, and leaves it equally good in all other respects, then the tweak makes for a better novel overall—if only barely.

Both parts of our definition are, admittedly, a bit fuzzy. But the basic ideas behind them should be uncontroversial. There is room to debate what counts as a "genuine and distinct" way of being good and about what counts as being "responsive" to changes in dimensions. And yet, we can agree that *having courage* may not be a genuine way of being a good person as distinct from *having the ideal amount of courage*, and nobody thinks that something can be good overall in virtue of being terrible in every dimension.

Moreover, the definition is useful even in its fuzzy form (indeed, we have used it several times already, and we will use it again). If imbalance between goods is in itself bad, then balance is a genuine and distinct way of being bad, which means that balance is a dimension of value. If more talent eventually makes for a worse philosopher, then betterness as a philosopher is not responsive in the right way to talent, which means that talent is not a dimension of philosophical value. If a tweak to a novel makes even a tiny improvement in all dimensions, then the result should be a novel that is overall better. More generally, if something is a Pareto improvement, then it ought to count as better overall; otherwise, value would not be properly responsive to improvements in dimensions (see §6 for more on this point).

But in addition to helping us analyze the cases used against Strong Pareto, our definition also sheds light elsewhere in ethics. For it reveals an ambiguity in the way we talk about "dimensions." Typically, the term refers to ways of being good. But sometimes, ethicists use 'dimension' to mean something more generic, like "a characteristic of the objected being evaluated." Temkin (2012: 304–5), for example, models the values of experiences with two dimensions: intensity and duration. For him, a value of, say, (1, 100) indicates a slight pleasure that lasts 100 hours, whereas (-100, 100) indicates an intense pain lasting for 100 hours. Is duration a dimension? Not in our sense: it measures how long an experience lasts, but that is not the same as *goodness with respect to duration*. The long intense pain is bad with respect to duration; the long dull pleasure is good with respect to duration; but their durations are the same.

These two senses of "dimension" need to be kept sharply distinct. Only the pluralist believes in multiple value "dimensions" in our sense—genuine and distinct ways of being good or bad to which overall value is responsive. But even monists can believe that the one true dimension has multiple characteristics. For example, a monist might see total pleasure as the only true dimension of value, while holding that pleasures can be characterized along the "dimensions" of duration and intensity. On this view, short intense pleasures are fungible with long dull ones: they are not good in distinct ways.

With that, we are done with the first two tasks of the paper. We have argued for a new definition of value dimensions (as genuine and distinct respects in which something can be good, to which overall value is responsive), and we have shown how that definition helps us defend the Pareto principles against a trio of counterexamples.

Our final task is to explore the idea of responsiveness more deeply, considering principles beyond Strong and Weak Pareto. Ideally, we would like to find principles that can serve as neutral constraints on axiology—strong enough to be interesting, yet still plausible given any reasonable theory. In our view, the two ideal principles for the job are Strong Pareto, which lays out the minimum way in which overall value must be responsive to dimensions (and which entails Weak

Pareto), and *Dimensionalism*, which adds that overall value is responsive *only* to dimensions (see §7).

A natural question, to start us off, is whether we should say something stronger than Strong Pareto. Can't we say more about how overall value responds to changes in dimensions? In the next section, we show how a common argument for Strong Pareto does indeed entail something stronger ("Super-Strong Pareto")—but also leads to a paradox. By contrast, our definition of dimensions supports Strong Pareto without generating the paradox; this is an additional virtue of our definition.

7 | THE "NO OBJECTION ARGUMENT" AND INCOMMENSURABILITY

Strong Pareto says that A is overall better than B if A is better in some way and at least as good in every other way. Although the principle is plausible in itself, and although we have defended it against various objections, we can also support it with arguments. The most familiar argument is as follows.

If A is better than B in some way, that counts in favor of A. If A is not worse in any other way, then there can be *no objection* to choosing A instead of B, and so nothing to defeat the point in favor of A. Since an undefeated point in favor is always decisive, A must emerge as overall better than B.

We call this the *No Objection Argument*. (The name is inspired by the argument's original form, where the dimensions of value correspond to the welfares of particular people, who cannot "object" when a change does not make them worse off (Murphy & Nagel, 2004 : 50; Adler, 2019: 96–7; Muñoz, 2022: 76).) If successful, the argument establishes Strong (and hence also Weak) Pareto.

This is a powerful argument—too powerful, we think. For if it works, it would also establish an ambitious principle that we call:

> **Super-Strong Pareto**
> If A is better than B along some dimension and not worse along any dimension, then A is overall better than B.[16]

This principle is equivalent to Strong Pareto if 'at least as good as' is equivalent to 'not worse than', as it is when things must be *commensurable* along all dimensions. To say that two things are commensurable (or "comparable") along a dimension is to say that they are either equally good along that dimension, or else one is better along it.

But what if things are *incommensurable* on a dimension—i.e., neither better, nor worse, nor equally good along that dimension?[17] Then Super-Strong Pareto will indeed be stronger. To illustrate, suppose a family trip to the beach is better with respect to price than a brunch with friends, but that the beach and brunch are incommensurable on every other dimension. Strong Pareto

[16] Thanks to a referee for suggesting this name, along with ways to trim this section.

[17] We use 'incommensurability' as an umbrella term to cover several related notions, like parity (Chang, 2002a, 2002b, 2005, 2016a, 2016b), imprecise equality (Parfit, 2016), rough equality (Griffin, 1986), and incomparability (RSinnott-Armstrong, 1985). For a guide to these concepts, see Andersson & Herlitz (2022).

does not entail that either is better overall, since neither is at least as good in every way.[18] But Super-Strong Pareto entails that the beach is better overall, since it is better in one way and no worse in any way (only incommensurable).

Now let's verify that the No Objection Argument would support Super-Strong Pareto. Sticking with the same case for illustration, the beach trip's better price counts in favor of its being better overall, and the rankings on other dimensions do not provide any objections, since they do not favor getting brunch instead; there is only incommensurability along these other dimensions. So, as before, the undefeated point in favor wins, and the beach trip is overall better.

But this reasoning must be mistaken, because Super-Strong Pareto is false. As Hare (2013: 177) notes, Super-Strong Pareto leads to *cycles*[19] of 'overall better than'.[20] To see this, suppose there are three options (A, B, C) and three dimensions of value: the welfares of three persons P1, P2, and P3. Suppose also that there can be incommensurability within each dimension. We represent this by letting values along each dimension be of two kinds, the x values (x, and x+) and the y values (y and y+), and by supposing that these x values are neither higher, nor lower, nor equal to these y values. Think of the x's as welfares of lives that are similar to each other, so that, e.g., x is the welfare of a certain life as a philosophy professor, y is the welfare of a certain life as a journalist, and x+ and y+ are the welfares of sweetened variants thereof: x+ is much like x yet slightly better, and y+ is much like y yet slightly better.

Now consider these values for our four options and three dimensions:

Dimension\Option	A	B	C
P1's welfare	y	x+	x
P2's welfare	x	y	x+
P3's welfare	x+	x	y

Here, Super-Strong Pareto says that A is better than B, for it is better with respect to P3's welfare and incommensurable (hence not worse) with respect to the welfares of P1 and P2. It says that B is better than C, for it is better with respect to P1's welfare and incommensurable with respect to the welfares of P2 and P3. But is also says that C is better than A, for it is better with respect to P2's welfare and incommensurable with respect to the welfares of P1 and P3. So, A is better than B, which is better than C, which is better than A: an intolerable betterness cycle![21]

[18] We assume that 'at least as good as' means better than or equally good'. Some, like Chang (2002a: 2), use it in a broader way that is compatible with incommensurability. Qizilbash (2021: 90) discusses one such "non-conventional" definition.

[19] x_1, x_2,…, x_n form a "cycle" of a relation R iff $x_1 R x_2$, $x_2 R x_3$,…, $x_{n-1} R x_n$, and $x_n R x_1$. If there exist cycles of R, and R is irreflexive or asymmetric, then R is nontransitive: $x_1 R x_2$ and $x_2 R x_3$ do not imply $x_1 R x_3$.

[20] Other theorists have been skeptical of Super-Strong Pareto (and related principles). Parfit (2016: 115) suggests that its false appeal comes from a simplistic "linear model" of values. Steele (2022) rejects a related principle that we call *Negative Pareto*: if A is not worse than B in any respect, then A is not worse overall. Steele links her principle to a deeper idea, which we call *Ignore Incommensurable Dimensions*: whether A is better overall than B depends only on the respects in which A and B are commensurable. Steele shows how this idea, and Negative Pareto, can lead to implausible conclusions. We are twisting Steele's knife, by showing that Strong Pareto, if we ignore incommensurable dimensions, leads to cycles of betterness.

[21] Some theorists may embrace this result and accept cyclic betterness (Rachels, 1998, 2001; Temkin, 1987, 2012), but most would not (Binmore & Voorhoeve, 2003; Broome, 1991; Huemer, 2013: 328). Given more options and more ways of sweetening, we can use Super-Strong Pareto to derive more exotic cycles than Hare's; we omit these for reasons of space—and because "ordinary" cycles are already strange enough.

Super-Strong Pareto, then, is false, even though it follows from the No Objection Argument. So what is wrong with that argument?

We think it is very plausible that A's being better than B along a dimension should favor A's being better overall. We also think that, if A and B are incommensurable along a dimension, that is no objection to choosing A. But incommensurability along a dimension may still count *against* either option's being overall better; it may instead count in favor of incommensurability overall.

There is, however, a core truth to the argument, which is really nothing more than our definition of dimensions—in particular, the second part of the definition, which says that overall value must be responsive to underlying dimensions. When A is better than B along a dimension, that counts in favor of A's being overall better, thereby ensuring that A will indeed be overall better *as long as nothing else counts against this*—in other words, as long as A is not worse than *or incommensurable with* B along any dimension.

It is this idea of responsiveness, we think, rather than the No Objection Argument, that supports Strong Pareto. Responsiveness tells us that when all else is equal (really *equal*!), the better something is along a given dimension, the better it is overall. Our definition of dimensions, and the idea of responsiveness therein, thus yields Strong and Weak Pareto, but it doesn't overgeneralize and yield Super-Strong Pareto. This tells us something about the real rationale for Strong Pareto, and it suggests that Strong Pareto and our definition are mutually supporting.

8 | DIMENSIONALISM AS A WORKING HYPOTHESIS

Strong Pareto, we have argued, follows from the fact that overall value is by definition responsive to dimensions of value. The minimal way to interpret this fact is that, other things being equal, a thing is improved when it is improved along a dimension.

In this final section, we focus on a different question about dimensions and responsiveness, on which our definition of dimensions is neutral. We've said that overall value is responsive to dimensions of value. But is that *all* that it is responsive to? Once we fix how A and B compare along each relevant dimension, do we thereby determine how A and B compare overall? We think the answer is *yes*. Overall value supervenes on values along underlying dimensions. Call this claim *Dimensionalism*.[22]

While we do not have a decisive argument for Dimensionalism, we think value theorists ought to adopt it as a working hypothesis. It has the appeal of simplicity and elegance. It's also plausible: *prima facie*, it seems that dimensions *just are* the relevant factors in overall comparisons. If something affects whether one thing is overall better than another, then it must affect how they compare along the relevant dimensions; if A is improved overall, it must have improved along a dimension.[23]

Despite its plausibility, however, Dimensionalism seems to be open to two kinds of counterexamples involving context-dependence.

[22] For a more extensive discussion and defense of Dimensionalism, see Hedden & Nebel (ms). As they note, Dimensionalism has an analogue in social choice theory: *Welfarism*, which holds that the values of states of affairs are solely determined by the welfare levels of individuals in those states; see Sen (1977). See Sen and Hedden & Nebel for discussion of how Dimensionalism/Welfarism can be imposed axiomatically on aggregation functions.

[23] Our thanks to a referee for helping us formulate this idea. Strictly speaking, this remark also relies on the background assumption that overall value can't be negatively responsive to dimensions.

First, on some views, context can make a dimension *irrelevant*. Imagine you are choosing which of two coats to buy for yourself: C$ is slightly better than C in every way, except that C is modestly better with respect to price. All else equal, C may be the overall better option. But now suppose you have promised a friend not to decide based on modest differences in price. Some ethicists might think that this promise makes the dimension of price irrelevant—Raz would call the promise an "exclusionary reason" (1999: 39–40).[24] If so, then the promise might make buying C$ the better option, even though it does not affect how good C$ is along any dimension; the promise merely makes one dimension (price) irrelevant. So it would seem that C$'s overall value can change even holding fixed each dimension.

But even granting that such context-relativity is possible, the Dimensionalist has an easy reply. When price becomes "irrelevant" (in a context), it ceases to be a relevant dimension of value (in that context).[25] Dimensionalism itself doesn't say that this sort of context-relativity is possible. But nor does it say the opposite! Dimensionalism just says: once we specify which dimensions are relevant in a given context, and how C and C$ compare along those dimensions, we thereby determine the overall values of C and C$.

Second, sometimes, context seems to affect the relative weights (rather than the existence) of dimensions. Suppose that a scientific theory's goodness depends on underlying dimensions that correspond to the "theoretical virtues" of simplicity, fit-with-the-data, fruitfulness, and so on (Kuhn, 1977). Morreau (2014: 1265) considers (without endorsing) the possibility that the relative weights of these theoretical virtues might vary with the domain of the theory in question. For instance, simplicity might matter more relative to fit for theories in physics than for theories in ecology. If so, then the overall goodness of a scientific theory would not depend solely on how it ranks on each underlying dimension, for it would also depend on the type of theory it is.

If a theory's domain affects its overall value, should the Dimensionalist just say that domain is *itself* a dimension? No, because this "dimension" is not a respect in which a theory can be good, to which overall value is responsive. That a theory is physical rather than ecological (or vice versa) does not *ipso facto* count in its favor.

Morreau's example is of exactly the right kind to challenge Dimensionalism. But we don't see it as very troubling, much less decisive. Even if we come to believe that theory domain can affect the relative weights of simplicity and fit, our first response should not be to give up on Dimensionalism. We should first ask: what if simplicity and fit are not the *fundamental* theoretical virtues? Is there some deeper reason why simplicity and fit play different roles in different domains of science? We can reasonably hope to find an answer that reveals some deeper dimensions of theoretical goodness, ones on which overall theoretical goodness supervenes.[26]

To be fair, we also see no a priori guarantee that this hope will always be borne out. Perhaps the search for more fundamental dimensions will be in vain. But we doubt it. Dimensionalism is immensely intuitive and surprisingly flexible. For now, we think any value pluralist in search of a systematic theory should take on Dimensionalism as a working hypothesis. If we

[24] See also Kamm (1993: 101; 2007: 34) on "irrelevant utilities" and Dancy (2004) on "disabling." Cullity (2013) provides an excellent overview of such contextual factors.

[25] For more on the question of which dimensions are relevant, see Chang (2004: 11).

[26] See Hedden & Nebel (ms) for this response to Morreau's case, as well as responses to related arguments from Hurley (1985, 1986: ch. 12).

ever do give up Dimensionalism, or abandon the search for systems, it should be as a last resort.

9 | CONCLUSION

In this paper, we have tried to show how value pluralists can be more systematic. First, we defended Strong and Weak Pareto—the bare minimum for any systematic pluralist—from a variety of counterexamples. This led us to a definition of dimensions of value: they are genuine and distinct respects in which something can be good or bad, to which overall value is responsive.

We then considered and rejected a stronger principle—Super-Strong Pareto—on the grounds that it yields cycles of betterness. The failure of Super-Strong Pareto tells us something about why Strong Pareto is true; it is true not on the basis of the fallacious No Objection Argument, but rather because overall value is responsive to dimensions. The upshot of responsiveness is that, holding all else equal (really *equal*!), improving something along one dimension should improve it overall.

We closed by exploring Dimensionalism, the thesis that overall value supervenes on which dimensions are relevant and how good things are along them. Without taking a definitive stand, we recommended that the pluralist adopt Dimensionalism as a working hypothesis. These are only some first steps toward a systematic pluralism. But we hope the rest of the journey will now appear more promising—or, at least, less perilous.[27]

ACKNOWLEDGEMENTS

Open access publishing facilitated by Australian National University, as part of the Wiley - Australian National University agreement via the Council of Australian University Librarians.

ORCID

Brian Hedden https://orcid.org/0000-0003-4463-9255
Daniel Muñoz https://orcid.org/0000-0001-6175-5675

WORKS CITED

Adler, M. (2019). *Measuring Social Welfare: An Introduction*. New York: Oxford University Press.

Andersson, H., & Herlitz, A. (2022). 'Introduction.' In H. Andersson & A. Herlitz (Eds.), *Value Incommensurability: Ethics, Risk, and Decision-Making* (pp. 1–25). New York: Routledge.

Aristotle. (1999). *Nicomachean Ethics*. Indianapolis: Hackett.

Binmore, K., & Voorhoeve, A. (2003). Defending Transitivity Against Zeno's Paradox. *Philosophy and Public Affairs*, *31*(3), 272–79. https://doi.org/10.1111/j.1088-4963.2003.00272.x

Broome, J. (1991). *Weighing Goods: Equality, Uncertainty and Time*. Cambridge, MA: Wiley-Blackwell.

Chang, R. (2002a). *Making Comparisons Count*. New York: Routledge.

Chang, R. (2002b). The Possibility of Parity. *Ethics*, *112*(4), 659–88. https://doi.org/10.1086/339673

Chang, R. (2004). 'All Things Considered.' *Philosophical Perspectives*, *18*(1), 1–22. https://doi.org/10.1111/j.1520-8583.2004.00018.x

Chang, R. (2005). Parity, Interval Value, and Choice. *Ethics*, *115*(2), 331–50. https://doi.org/10.1086/426307

[27] Our sincere thanks to the editors at Noûs, as well as to our anonymous referee, who suggested ways to streamline the paper (e.g., in §§5.6), shore up the arguments (especially in §3), and cut down on symbols (throughout). We would also like to thank Kerah Gordon-Solmon and the students in her Ethics seminar at Queen's University.

Chang, R. (2016a). Parity: An Intuitive Case. *Ratio*, *29*(4), 395–411. https://doi.org/10.1111/rati.12148
Chang, R. (2016b). Parity, Imprecise Comparability, and the Repugnant Conclusion. *Theoria*, *82*(2), 183–215. https://doi.org/10.1111/theo.12096
Chappell, R. Y. (2015). Value Receptacles. *Noûs*, *49*(2), 322–32. https://doi.org/10.1111/nous.12023
Constantinescu, C. (2012). Value Incomparability and Indeterminacy. *Ethical Theory and Moral Practice*, *15*(1), 57–70. https://doi.org/10.1007/s10677-011-9269-8
Cullity, G. (2013). The Context-Undermining of Practical Reasons. *Ethics*, *124*(1), 8–34. https://doi.org/10.1086/671390
Dancy, J. (2004). *Ethics Without Principles*. Oxford: Oxford University Press.
De Sousa, R. (1974). The Good and the True. *Mind*, *83*(332), 534–51.
Gaertner, W. (2009). *A Primer in Social Choice Theory: Revised Edition*. Oxford: Oxford University Press.
Griffin, J. (1986). *Well-Being: Its Meaning, Measurement and Moral Importance*. Oxford: Clarendon Press.
Hare, C. (2013). *The Limits of Kindness*. Oxford: Oxford University Press.
Hedden, B., & Nebel, J. (ms). Multidimensional Concepts and Disparate Scale Types.
Hsieh, N. (2005). Equality, Clumpiness and Incomparability. *Utilitas*, *17*(2), 180–204. https://doi.org/10.1017/S0953820805001512
Huemer, M. (2013). Transitivity, Comparative Value, and the Methods of Ethics. *Ethics*, *123*(2), 318–45. https://doi.org/10.1086/668905
Hurka, T. (1993). *Perfectionism*. New York: Oxford University Press.
Hurley, S. (1985). Supervenience and the Possibility of Coherence. *Mind*, *94*(376), 501–25. https://doi.org/10.1093/mind/XCIV.376.501
Hurley, S. (1989). *Natural Reasons: Personality and Polity*. Oxford: Oxford University Press.
Kamm, F. (1993). *Morality, Mortality: Volume I: Death and Whom to Save From It*. New York: Oxford University Press.
Kamm, F. (2007). *Intricate Ethics: Rights, Responsibilities, and Permissible Harm*. New York: Oxford University Press.
Kuhn, T. (1977). "Objectivity, Value Judgment, and Theory Choice." In T. Kuhn (Ed.), *The Essential Tension: Selected Studies in Scientific Tradition and Change* (pp. 320–39). Chicago: University of Chicago Press.
Morreau, M. (2014). Mr. Fit, Mr. Simplicity, and Mr. Scope: From Social Choice to Theory Choice. *Erkenntnis*, *79*, 1253–68. https://doi.org/10.1007/s10670-013-9549-x
Muñoz, D. (forthcoming). Sources of Transitivity. Economics and Philosophy.
Muñoz, D. (2022). The Many, the Few, and the Nature of Value. *Ergo*, *9*(4), 70–87. https://doi.org/10.3998/ergo.2260
Murphy, L., & Nagel, T. (2004). *The Myth of Ownership: Taxes and Justice*. Oxford: Oxford University Press.
Nagel, T. (1979). The Fragmentation of Value. In T. Nagel (Ed.), *Mortal Questions* (pp. 128–41). Cambridge: Cambridge University Press.
Nebel, J. (2015). *Incommensurability in Population Ethics*. Unpublished BPhil dissertation.
Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford University Press.
Parfit, D. (1997). Equality and Priority. *Ratio*, *10*(3), 202–21. https://doi.org/10.1111/1467-9329.00041
Parfit, D. (2016). Can We Avoid the Repugnant Conclusion? *Theoria*, *82*(2), 110–27. https://doi.org/10.1111/theo.12097
Qizilbash, M. (2021). On 'Incommensurability,' 'Discontinuity,' and the Repugnant Conclusion: 'Imprecise Equality' or Vagueness? In H. Andersson & A. Herlitz (Eds.), *Value Incommensurability: Ethics, Risk, and Decision-Making* (pp. 87–108). New York: Routledge.
Rachels, S. (1998). Counterexamples to the Transitivity of Better Than. *Australasian Journal of Philosophy*, *76*(1), 71–83. https://doi.org/10.1080/00048409812348201
Rachels, S. (2001). A Set of Solutions to Parfit's Problems. *Noûs*, *35*(2), 214–38. https://doi.org/10.1111/0029-4624.00294
Raz, J. (1999). *Practical Reason and Norms*. Oxford: Oxford University Press.
Ross, W. D. (1930). *The Right and the Good*. Oxford: Oxford University Press.
Sen, A. (1977). On Weights and Measures: Informational Constraints in Social Welfare Analysis. *Econometrica*, *45*(7), 1539–72. https://doi.org/10.2307/1913949
Sen, A. (2017). *Collective Choice and Social Welfare: An Expanded Edition*. Cambridge: Harvard University Press.
Sinnott-Armstrong, W. (1985). Moral Dilemmas and Incomparability. *American Philosophical Quarterly*, *22*(4), 321–29.

Steele, K. (2022). Incommensurability That Can(Not) Be Ignored. In H. Andersson & A. Herlitz (Eds.), *Value Incommensurability* (pp. 231–46). New York: Routledge.

Temkin, L. (1983). *Inequality*. PhD Dissertation, Princeton University.

Temkin, L. (1987). Intransitivity and the Mere Addition Paradox. *Philosophy and Public Affairs*, *16*(2), 138–87.

Temkin, L. (2012). *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*. Oxford: Oxford University Press.

Wolf, S. (1992). Two Levels of Pluralism. *Ethics*, *102*(4), 785–98. https://doi.org/10.1086/293448

How to cite this article: Hedden, B., & Muñoz, D. (2023). Dimensions of Value. *Noûs*, 1–15. https://doi.org/10.1111/nous.12454