

Multidimensional Concepts and Disparate Scale Types

Brian Hedden and Jacob M. Nebel

Abstract

Multidimensional concepts are everywhere, and they are important. Examples include moral value, welfare, scientific confirmation, democracy, and biodiversity. How, if at all, can we aggregate the underlying dimensions of a multidimensional concept F to yield verdicts about which things are F er than which *overall*? Social choice theory can be used to model and investigate this aggregation problem. Here, we focus on a particularly thorny problem made salient by this social choice-theoretic framework: the underlying dimensions of a given concept might be measurable on different types of scales—e.g., some ordinal and some cardinal. An underappreciated impossibility theorem due to Anna Khmelnitskaya shows that seemingly plausible constraints on aggregation across scale types are inconsistent. This impossibility threatens to render the notion of overall F ness incoherent. We attempt to defuse this threat, arguing that the impossibility depends on an overly restrictive conception of measurement and of how measurement constrains aggregation. Adopting a more flexible—and, we think, more perspicuous—conception of measurement opens an array of possibilities for aggregation across disparate scale types.

1 Introduction

Many of our most important concepts are multidimensional. A concept F is multidimensional just in case whether and to what extent something is F depends on

For helpful feedback and discussion, we would like to thank Christopher Bottomley, Kevin Dorst, Keith Dowding, Bob Goodin, Zach Goodsell, Alan Hájek, Harvey Lederman, Daniel Muñoz, John Weymark, two anonymous referees, and an audience at the Australian National University.

how it stands along multiple underlying dimensions, or respects, of *Fness*. Take *democratic*. Whether and to what extent a country is democratic depends not only on whether it holds elections, but also on voter turnout rates, the protection of civil liberties, independence of the judiciary, and so on. Or take *good*. According to pluralists about value, whether and to what extent a state of affairs is good depends not only on how much welfare it contains, but also on how it stands with respect to other values such as equality and freedom.

How do multidimensional concepts work? How, in particular, can the underlying dimensions of a multidimensional concept be aggregated to yield overall, or all things considered, judgments and comparisons? This question can be addressed using tools from social choice theory. Social choice theory studies how—if at all—individual preferences or utilities can be aggregated to yield an overall “social” ordering of alternatives. But we can apply the framework of social choice theory to multidimensionality more generally; the key is to think of each underlying dimension of a multidimensional concept as akin to an individual whose preferences or utilities correspond to that dimension’s ranking of alternatives.

A problem immediately emerges: In the traditional setting of social choice theory, it is plausible that all individuals’ utilities are measurable on the same *type* of scale—e.g., ordinal, interval, or ratio—though there is controversy about what type that is. By contrast, there is no reason to assume that all underlying dimensions of a multidimensional concept will be measurable on the same type of scale. Instead, many multidimensional concepts are likely to have dimensions measurable on disparate scale types. How, if at all, can we aggregate across scale types?

There are no obvious answers. And while this problem has received scant attention among social choice theorists, the only work on the problem is pessimistic. An impossibility theorem due to Anna Khmelnitskaya (1996, 2002, 2010) (see also Khmelnitskaya and Weymark, 2000) shows that plausible constraints on ways of ranking alternatives using information from disparate scale types are jointly unsatisfiable. This result might be taken to show that, for many apparently multidimensional concepts *F*, there is no such thing as overall *Fness*. Something can be *F along some dimension* or *in some respect*, but we cannot aggregate these dimensions or respects so as to yield verdicts about whether and to what extent it is *F overall* or *all things considered*. Big news, if true!

In this paper, we attempt to tackle this thorny and important problem. We begin

by setting out the problem in more detail, describing multidimensionality and the applicability of social choice theory thereto and outlining different types of measurement scales. Then, after stating Khmelnitskaya's impossibility theorem, we consider a variety of potential "escape routes." Many of these strategies, we argue, are not sufficiently general to deal with the problem of disparate scale types in its entirety and in a plausible way. However, one response to the problem does seem to fit the bill, namely that the impossibility theorem rests on an overly restrictive conception of measurement and of how measurement constrains aggregation.

To clarify, our aim in this paper is not to propose or defend any particular way of making multidimensional comparisons with disparate scale types. We suspect that there is no one-size-fits-all method to be discovered. Rather, our aim is to defuse a particular threat to the possibility of such comparisons, based on Khmelnitskaya's impossibility theorem. Our particular solution has the virtue of being compatible with a very wide range of methods for making multidimensional comparisons. It is a task for future research to narrow that range down to classes of plausible methods which are especially well suited to particular concepts.

2 Multidimensional Concepts and Aggregation

Multidimensionality is ubiquitous. Consider some further examples, beyond those of democracy and value. Biodiversity is widely held to be multidimensional, depending not only on the number of species present in the ecosystem, but also their phylogenetic and morphological diversity (Maclaurin and Sterelny, 2008). Inequality is multidimensional, depending on inequality of income, inequality of resources, and inequality of opportunity, among other things, each of which may have multiple aspects (Sen, 1997; Temkin, 1993). Welfare is widely held to be multidimensional: objective list theorists (or pluralists) hold that welfare consists not only in happiness or pleasure, but also in preference satisfaction, love, knowledge, and so on (Lin, 2014). Similarly for scientific confirmation: the extent to which a theory is credible depends on how it stands with respect to a plurality of theoretical virtues like simplicity, fit with the data, scope, and fruitfulness (Kuhn, 1977). Under normative uncertainty, overall choiceworthiness seems multidimensional, depending on how an action ranks relative to others according to the various first-order normative theories that one takes seriously (Sepielli, 2009). Overall similarity—a

central part of Lewis (1973)'s theory of counterfactuals—is likewise multidimensional, depending on a number of underlying respects of similarity. And in the context of a normality-theoretic approach to epistemology, Goodman and Salow (2023) suggest that normality is multidimensional, as situations can be more or less normal in different respects.

All of the multidimensional concepts just mentioned are gradable; they admit of a comparative form. But there are also some non-gradable concepts which are arguably multidimensional. For Lewis (1973), the concept of a law of nature is multidimensional, for the laws are the axioms or theorems in the best true deductive system for capturing facts about the “Humean mosaic,” where the best true deductive system is the one which strikes the best balance of simplicity and informational strength. For similar reasons, he thinks that the concepts of belief and desire are multidimensional; roughly speaking, an agent has the beliefs and desires whose attribution would give the best interpretation of their behavior, where the best interpretation is the one which strikes the best balance of charity and fit with behavior (Lewis, 1974). Analogously, for Dworkin (1986), the concept of being the correct legal interpretation is multidimensional: the correct interpretation is the one which strikes the best balance between charity (casting the law in the best moral light possible) and fit with text and precedent. But these concepts are non-gradable; nothing is “law-ier” or “belief-ier” than anything else. For ease of exposition, we’ll usually talk about multidimensional gradable concepts in what follows, but what we say will typically apply to non-gradable ones as well—especially since many accounts of non-gradable multidimensional concepts, like those just mentioned, analyze them in terms of gradable ones (e.g., the *best* system or interpretation), and we can often construct complex comparative expressions (e.g., *more lawlike*) out of the non-gradable ones.

Sometimes, we may be content to “disaggregate” and restrict ourselves to talking about whether and to what extent something is *F along a given dimension* or *in a given respect*. But often, we want to make overall judgments and comparisons. We want to talk about which countries are overall more democratic than which, whether overall biodiversity is decreasing, which policy is overall best, and so on. Similarly, overall comparisons are required in order for Lewis’s theories to yield verdicts about what the laws of nature are, what beliefs and desires an agent has, and which counterfactuals are true. They are also required for Dworkin’s theory to

yield verdicts about the proper interpretation of the law and for normality-theoretic approaches to epistemology to yield verdicts about what agents know.

Such overall judgments and comparisons require that the underlying dimensions of a multidimensional concept be somehow aggregated. How, if at all, can this be done? Here, controversy reigns. In ethics, some theorists think that the multidimensionality of value means that the *at least as good as* relation is intransitive (Rachels, 1998; Temkin, 2012), incomplete (Raz, 1985; Chang, 2002b), or both.¹ In epistemology, some permissivists hold that there is no uniquely privileged way of aggregating the competing theoretical virtues, and so there is often no uniquely privileged doxastic state, given a body of evidence (Schoenfield, 2014). Goodman (1972) doubted the coherence of overall similarity.

We can use social choice theory to investigate whether and how the dimensions of a multidimensional concept can be aggregated to yield overall judgments and comparisons. In the traditional setting of welfare economics, social choice theory is concerned with how, if at all, individual preferences or utilities can be aggregated to yield a single overall “social” ordering. We seek an aggregation function (often called a *social welfare function*) that takes as input a list of either preference orderings or utility functions, one per individual, and outputs a single overall ordering.

This mathematical framework can be interpreted so as to apply to multidimensional concepts more generally. Just think of each underlying dimension of a multidimensional concept as corresponding to an individual whose preferences or utilities match the ranking of alternatives along that dimension. In this setting, we seek an aggregation function that takes as input a list of orderings or utility functions, one per underlying dimension, and outputs a single overall ordering.

Others have used social choice theory to model particular multidimensional concepts of philosophical interest. Hurley (1985) has done so for value pluralism, Morreau (2010) and Kroedel and Huber (2013) for overall similarity, Okasha (2011) for scientific confirmation, MacAskill et al. (2020) for normative uncertainty, and Hattiangadi (2020) for interpretivism about the mind.² List (2002, 2004) considers the problem of aggregating simultaneously across individuals and dimensions of

¹A binary relation R on a set X is transitive just in case for all x, y, z in X , if xRy and yRz , then xRz . And it is complete just in case for all x, y in X , either xRy or yRx (or both).

²See also Grinsell (2012, 2017); van Rooij (2011); D’Ambrosio and Hedden (forthcoming) for relevant work in linguistics and philosophy of language that applies social choice theory to the semantics of multidimensional adjectives.

welfare; he even makes use of Khmelnitskaya's results. But we are concerned with multidimensional concepts in full generality, and with a specific problem that arises in analyzing them through the lens of social choice theory, a problem to which we now turn.

3 Measurability and Disparate Scale Types

In the traditional welfare economic setting of social choice theory, there has been much debate about the type of scale on which individual preferences or welfares are measurable. Arrow (1951) held that individual welfare, understood in terms of individual preferences, is merely ordinally measurable (and, moreover, interpersonally noncomparable). Alternatives can be better or worse for each person, but not by more or less. In this case, our aggregation function can take as input a list of individual preference *orderings*. We could have it take as input a list of real-valued *utility functions*, but we would need to be sensitive to the fact that if some individual's preferences are merely ordinally measurable and some real-valued utility function u represents those preferences, then so does any strictly increasing (i.e. order-preserving) transformation thereof—i.e., any u^* such that $u^*(x) > u^*(y)$ iff $u(x) > u(y)$.

Later theorists (Sen, 1970a; d'Aspremont and Gevers, 1977; Roberts, 1980a) suggested that individual welfare may be informationally richer than Arrow assumed. In particular, welfare may be measurable on an *interval scale*. If so, then one thing can be better for a person than another by more or less than a third thing is better than a fourth. More specifically, an interval scale represents meaningful *ratios of differences*, like the Celsius and Fahrenheit scales of temperature. But there are no meaningful welfare *ratios* on such a scale—e.g., something being twice as good for a person as another—for this would require a privileged origin, or zero point, for welfare scales. Interval scales have meaningful cardinal structure but an arbitrary unit and origin: if utility function u represents some individual's welfare, then so does any positive affine transformation thereof—i.e., any $u^* = au + b$, where $a > 0$.

Finally, theorists such as Broome (2004) and Adler (2011) have proposed that individual welfare is still more informationally rich. In particular, it is measurable on a *ratio scale*, like mass and length. A ratio scale has a meaningful origin but an arbitrary unit. So if welfare is ratio-scale measurable and if utility function u

represents some individual's welfare, then so does any similarity transformation thereof—i.e., any $u^* = au$, where $a > 0$. Any such transformation preserves ratios of utilities, and thus ratios of the welfare levels they represent (e.g., something being twice as good for you).

Here's the kicker: in the context of welfare economics, while there is much dispute about the type of scale on which individual welfare is measurable, it is nonetheless plausible that each individual's welfare is measurable on the same type of scale, whatever that may be. By contrast, in the context of multidimensional concepts generally, there is *prima facie* no reason to think that all underlying dimensions of a given concept will be measurable on the same type of scale. Instead, some might be ordinal-scale measurable, others interval-scale measurable, and still others ratio-scale measurable.³ Indeed, there is reason to think that at least some multidimensional concepts will have underlying dimensions measurable on disparate scale types, given the sheer number and diversity of such concepts.

Take, for example, the concept of welfare, with underlying dimensions of pleasure, preference satisfaction, and contemplation of beauty. We might think that amounts of pleasure are measurable on a ratio scale, while degrees of preference satisfaction are measurable on an interval scale, and contemplation of beauty only on an ordinal scale. Or take scientific confirmation, with underlying dimensions of simplicity, fit with the data, and fruitfulness. Perhaps simplicity and fit with the data are ratio-scale measurable, while fruitfulness is merely ordinally measurable (Okasha, 2011, p. 103). Or take overall choiceworthiness in the context of normative uncertainty. Some first-order moral theories (e.g., some deontological theories) seem only to deliver ordinal rankings of options, while others represent the value of options on an interval scale (e.g., decision-theoretic consequentialism), and others still may supply a ratio scale (Hedden, 2016; MacAskill et al., 2020; Tarsney, 2021). Of course, no such example is incontestable. But again, the sheer number and diversity of multidimensional concepts suggests that the problem of disparate scale types will arise for at least some of them.

Our question, then, is this: How—if at all—can we aggregate underlying dimensions measurable on disparate scale types to yield an overall, or all-things-considered, ordering?

³There are other possible scale types, too, such as translation scales, log-interval scales, absolute scales, Sen (1970a)'s uncountable spectrum of 'ordinal-type' scales, and Luce et al. (2014, ch. 20)'s scales of intermediate strength between interval and ratio types; we'll ignore this added complexity.

There are no off-the-shelf, well-studied aggregation functions for aggregating dimensions of disparate scale types. Offhand, it seems we would need to either “dumb down” the more informationally rich dimensions (e.g., treating interval-scale measurable dimensions as merely ordinal) or else “smart up” the more informationally impoverished dimensions (e.g., representing each ordinal-scale measurable dimension by its “Borda score” of how many alternatives fare worse along that dimension).⁴ Both are *prima facie* problematic. Dumbing down involves ignoring potentially useful and important information, while smarting up involves arbitrariness. Worse still, while there has been scant attention to the problem of disparate scale types by social choice theorists, the sole work we know of is pessimistic. We turn to that now.

4 Khmelnitskaya’s Impossibility Theorem

We begin with some setup. We have a set of at least three objects $X = \{x, y, z, \dots\}$, which we generically call “alternatives.” We want to know which alternatives are overall at least as F as which. Suppose that the F ness of an alternative depends on how it fares along n dimensions in some set $N = \{1, 2, \dots, n\}$. Each dimension i has a dimensional *utility function* u_i , which assigns a real number to each alternative in X . This implies that each dimension is measurable on some real-valued scale, though it does not settle the *type* of scale (i.e., ordinal, interval, or ratio).

A *profile* $U = (u_1, \dots, u_n)$ is an n -tuple of utility functions, one for each dimension of F ness. An *aggregation function* assigns, to each profile U in its domain, a transitive and complete at-least-as- F -as ordering on X . The aggregation function tells us, given any way in which the alternatives might be evaluated by the various dimensions, which of those alternatives are at least as F as which. Since the ordering delivered by the aggregation function depends on the profile taken as input, we index F to the profile U , as in ‘at least as F_U ’, ‘more F_U ’ and ‘equally F_U .’ We tell the aggregation function how the alternatives compare according to each dimensional utility function, and it tells us how the alternatives compare overall with respect to F given the values assigned by the underlying dimensions.

Khmelnitskaya assumes that the aggregation function looks only at the values

⁴See Tarsney (2021), who discusses “structural enrichment” and “structural depletion” in the context of decision-making under normative uncertainty.

assigned by the various dimensional utility functions. In social choice theory, this assumption is called *Welfarism*, because it means that the ranking of alternatives depends only on individual welfare levels. In our more general setting, we call this assumption *Dimensionalism*. Dimensionalism can be secured by imposing two conditions:

Pareto Indifference For any profile $U = (u_1, \dots, u_n)$ and alternatives $x, y \in X$, if $u_i(x) = u_i(y)$ for every dimension $i \in N$, then x and y are equally F_U .

Independence of Irrelevant Alternatives For any profiles $U = (u_1, \dots, u_n)$ and $U' = (u'_1, \dots, u'_n)$ and alternatives $x, y \in X$, if $u_i(x) = u'_i(x)$ and $u_i(y) = u'_i(y)$ for every dimension $i \in N$, then x is at least as F_U as y iff x is at least as $F_{U'}$ as y .

The combined force of these conditions depends on the domain of the aggregation function. Khmelnitskaya assumes it to be *unrestricted*, or defined for all profiles of utility functions:

Unrestricted Domain For any n -tuple of utility functions $U = (u_1, \dots, u_n)$, there exists a corresponding at-least-as- F_U -as ordering.

Pareto Indifference and Independence of Irrelevant Alternatives greatly simplify the workings of an aggregation function. They allow us to abstract from the alternatives and the profiles and just focus on the dimensional utilities that might be assigned. More specifically, given a profile $U = (u_1, \dots, u_n)$, an alternative x 's *utility vector* is the list of numbers that each dimension's utility function assigns to x : $U(x) = (u_1(x), \dots, u_n(x))$. The Welfarism (or Dimensionalism) Theorem says that, given Unrestricted Domain, Pareto Indifference and Independence of Irrelevant Alternatives are together equivalent to the existence of a single ordering of utility vectors which determines the ordering of alternatives assigned to any given profile (Bossert and Weymark, 2004, Theorem 2.2). This allows us to characterize some utility vectors as being at least as F as others, without indexing to a profile. (Strictly speaking, this is a category mistake, since utility vectors are just lists of numbers, and the multidimensional concepts we are interested in don't apply to such objects. But we treat it as shorthand for "anything with this utility vector is at least as F as anything with that one.")

In order to capture the idea of different scale types, Khmelnitskaya uses the framework of *invariance conditions* developed by Sen (1970a). In this framework, two profiles are deemed *informationally equivalent* when they represent the same meaningful information. Which profiles are informationally equivalent depends on the measurability and comparability of the various dimensions. The aggregation function is then required to assign the same at-least-as- F -as ordering to profiles that are informationally equivalent (see Weymark 2016 for a helpful overview).

For example, suppose we have two utility profiles $U = (u_1, \dots, u_n)$ and $\phi(U) = (\phi_1(u_1), \dots, \phi_n(u_n))$. If all underlying dimensions are ordinally measurable and we cannot make comparisons across distinct dimensions, then the two profiles would be said to be informationally equivalent just in case each ϕ_i is a strictly increasing transformation, possibly a different one for each i . Thus the profiles would have to be assigned the same ordering by the aggregation function. The idea, more generally, is that the numerical representation of the various dimensions' orderings of the alternatives is unique only up to certain classes of transformations, so profiles related by such transformations are informationally equivalent and therefore must be assigned the same at-least-as- F -as ordering. Given the Dimensionalism (or Welfareism) axioms of Pareto Indifference and Independence of Irrelevant Alternatives, this implies that pairs of utility vectors related by such transformations ought to be compared the same way.

Khmelnitskaya's key innovation is to consider invariance conditions that allow different groups of people (dimensions) to differ in terms of the measurability and interdimensional comparability of their members. Suppose that the n dimensions can be partitioned into m (≥ 2) subgroups, where all of the dimensions within each subgroup belong to the same scale type; crucially, dimensions in distinct subgroups are not comparable to one another.⁵ For example, suppose that dimensions $1, \dots, k$ are measurable on an ordinal scale that allows for interdimensional comparisons, whereas dimensions $k + 1, \dots, n$ are measurable on an interval scale without interdimensional comparability. Then profiles U and $\phi(U)$ would be considered informationally equivalent just in case, for every dimension i in the first subgroup, $\phi_i = \phi_0$ for some strictly increasing transformation ϕ_0 , and for every dimension j in the second subgroup, ϕ_j is a positive affine transformation—i.e., $\phi_j(u_j) = a_j u_j + b_j$

⁵Note that Khmelnitskaya's use of (non-singleton) subgroups is needed only if there is some interdimensional comparability between certain dimensions. If no dimensions are comparable, then we can equivalently treat each dimension as belonging to its own singleton subgroup.

($a_j > 0$), with a_j and b_j being possibly different for each j .⁶

We can state these invariance conditions more precisely as follows. Let Γ be the set of subgroups. For any utility profile U and subgroup $G \in \Gamma$, let U_G be the subprofile of utility functions for all and only those dimensions in G . For each subgroup G , there is a set Φ_G of *subgroup invariance transforms*. Each transform $\phi_G \in \Phi_G$ takes a subprofile U_G and returns a new one $\phi_G(U_G)$, where each G -dimension's utility function u_i is mapped to $\phi_i(u_i)$. The thought is that subprofiles related by transforms in Φ_G represent the same meaningful information about how the dimensions in G evaluate the various alternatives.

The version of Khmelnitskaya's impossibility theorem that we are interested in requires each subgroup to have one of seven types of scales. We have already mentioned three, along with their associated invariance conditions: ordinal measurability without interdimensional comparability, ordinal measurability with interdimensional comparability, and interval-scale measurability without interdimensional comparability. Interval-scale measurability with full interdimensional comparability requires the positive affine transformation to be the same for each dimension (within the relevant subgroup). Interval-scale measurability with interdimensional *unit* comparability requires the positive scale factor (a_j above) for each dimension to be the same (within the relevant subgroup), but it allows the translation (b_j) to differ by dimension. Ratio-scale measurability without interdimensional comparability requires each transformation to be a similarity transformation, possibly a different one for each dimension (within the relevant subgroup). And ratio-scale measurability with full interdimensional comparability requires that each transformation be the same similarity transformation for each dimension. Thus, we have the following possibilities for subgroup invariance transforms:

Ordinal Scales without Comparability (ONC) $\phi_G \in \Phi_G$ iff, for all $i \in G$, ϕ_i is strictly increasing.

Ordinal Scales with Comparability (OFC) $\phi_G \in \Phi_G$ iff there is some strictly increasing transformation ϕ_0 such that, for all $i \in G$, $\phi_i = \phi_0$.

Interval Scales without Comparability (INC) $\phi_G \in \Phi_G$ iff, for all $i \in G$, ϕ_i is a positive affine transformation.

⁶A related problem is studied by Accinelli and Plata (2008) who, building on Sen (1970b), consider individuals with partially comparable interval scales of various kinds.

Interval Scales with Full Comparability (IFC) $\phi_G \in \Phi_G$ iff there is some positive affine transformation ϕ_0 such that, for all $i \in G$, $\phi_i = \phi_0$.

Interval Scales with Unit Comparability (IUC) $\phi_G \in \Phi_G$ iff there is some positive real number a such that, for all $i \in G$, ϕ_i is a positive affine transformation with scale factor a .

Ratio Scales without Comparability (RNC) $\phi_G \in \Phi_G$ iff, for all $i \in G$, ϕ_i is a similarity transformation.

Ratio Scales with Full Comparability (RFC) $\phi_G \in \Phi_G$ iff there is some similarity transformation ϕ_0 such that, for all $i \in G$, $\phi_i = \phi_0$.

Khmelnitskaya requires the aggregation function to assign the same ordering to any two profiles related by a suitable transformation of each subgroup's utility functions corresponding to its scale type, where each subgroup's scale type is one of the seven just mentioned. The appropriate classes of transformations are mutually independent, since we are allowing interdimensional comparability only within subgroups, and not across them. Formally:

Informational Invariance across Subgroups For any profiles U and V , if for every subgroup $G \in \Gamma$, there is some $\phi_G \in \Phi_G$ such that $V_G = \phi_G(U_G)$, then: for any alternatives $x, y \in X$, x is at least as F_U as y iff x is at least as F_V as y .

Khmelnitskaya's next two conditions are easiest to state in terms of the ordering of utility vectors. The first says that having a greater value along every dimension is sufficient for being *Fer* overall:

Weak Pareto For any utility vectors u and v , if $u_i > v_i$ for every dimension $i \in N$, then u is *Fer* than v .

The second prohibits the ordering of utility vectors from being hypersensitive to arbitrarily small changes in the values assigned by various dimensions. More formally, say that a *neighborhood* Z_u of a utility vector u is a set containing all vectors that are within a certain (Euclidean) distance of u . According to

Continuity For any utility vectors u and v , if u is *Fer* than v , then there are neighborhoods Z_u and Z_v of u and v such that every u' in Z_u is *Fer* than every v' in Z_v .

Finally, let us say that a subgroup G is *strongly dictatorial* if the at-least-as- F -as ordering of alternatives depends *only* on the utilities assigned to those alternatives by dimensions within subgroup G . Formally, this means that there is an at-least-as- F -as ordering on the set of all $|G|$ -vectors such that, for any profile U and alternatives x and y , x is at least as F_U as y iff $U_G(x)$ is at least as F as $U_G(y)$. Hence, if there is a strongly dictatorial subgroup, all dimensions outside this subgroup are simply ignored, *always*. The final condition prohibits such a subgroup:

Weak Subgroup Nondictatorship There is no strongly dictatorial subgroup.

We can now state the result with which we are concerned:

Khmelnitskaya’s Impossibility Theorem. *Suppose the set of dimensions N can be partitioned into $m \geq 2$ subgroups $\Gamma = \{G_1, \dots, G_m\}$, each of which has one of the following seven scale types: ONC , OFC , INC , IUC , IFC , RNC , RFC . Then there is no aggregation function which satisfies Unrestricted Domain, Pareto Indifference, Independence of Irrelevant Alternatives, Weak Pareto, Continuity, Informational Invariance across Subgroups, and Weak Subgroup Nondictatorship.*

This particular result is not stated explicitly in any of Khmelnitskaya’s papers. But it follows straightforwardly from some of her results (Khmelnitskaya, 2002, Theorems 3.1 and 3.5). We give a simpler proof in Appendix A. (Our proof extends a strategy used by Khmelnitskaya and Weymark 2000 to establish a weaker version of the theorem involving only ONC , OFC , INC , IUC , and IFC subgroups.)

Khmelnitskaya’s impossibility theorem says that there must be a strongly dictatorial subgroup if the aggregation function satisfies the other conditions of the theorem. This is an incredible result. The conditions that give rise to a strongly dictatorial subgroup seem applicable to at least some multidimensional concepts, but the existence of such a subgroup seems applicable to none of them. After all, in what sense could some factor be a *dimension* of F_{ness} if it is completely irrelevant to the ordering of alternatives with respect to F_{ness} ? This is why we regard Khmelnitskaya’s result as a challenge for multidimensional concepts. (Things get even worse if the strongly dictatorial subgroup is of type ONC , INC , or RNC . Then, exactly *one* dimension within that subgroup fully determines the ordering assigned to each profile: that is, there is a strongly dictatorial dimension, not just a strongly dictatorial subgroup. This is because, for purposes of Informational Invariance across

Subgroups, there is no formal difference between a subgroup of k noncomparable dimensions and k singleton subgroups.)

In the remainder of this paper, we consider ways to avoid this pessimistic impossibility. We won't try to survey all possible escape routes, instead restricting ourselves to three strategies that strike us as particularly intriguing, and where we think we have novel contributions to make. In section 5, we consider the possibility of rejecting Continuity; in section 6, we consider the possibility of rejecting Dimensionalism; and in section 7, we consider the possibility of rejecting the informational invariance framework in general, and Informational Invariance across Subgroups in particular. We regard this last strategy as most promising.

Before turning to these strategies, however, we close this section by briefly flagging four others that we regard as less promising. First, we might reject the demand that aggregation functions output an at-least-as- F -as *ordering*, which must be both transitive and complete. It is easy to see that the remaining conditions can be satisfied once we jettison transitivity, completeness, or both. Rejecting transitivity opens space for Simple Majority Rule, for example, whereby x is overall at least as F as y just in case x is at least as F as y on a majority of underlying dimensions. And rejecting completeness opens space for the Strong Pareto Rule, whereby x is at least as F as y overall just in case x is at least as F as y on all underlying dimensions.⁷ There is, of course, extensive debate about both transitivity and completeness: Rachels (1998) and Temkin (2012) famously reject transitivity, while Raz (1985) and Chang (2002b) reject completeness. But we are sympathetic to the orthodox view on which transitivity and completeness hold for all comparatives, including multidimensional ones. We won't attempt to defend this orthodoxy here, but see Broome (2004) and Nebel (2018) for defenses of transitivity and Dorr et al. (2023) for a defense of completeness.

Second, and perhaps more radically, we might reject the demand that aggregation functions output a full ranking of alternatives, settling instead for ones which select a "winner" (or a set of winners) from any given menu of alternatives (when given a utility profile as input). In the jargon, we might settle for a *functional collective choice rule* (FCCR), rather than a *social welfare functional*. Adopting this more

⁷These are particularly simple transitivity- and completeness-violating aggregation functions. There are various more complex aggregation functions that violate these constraints, like the lexicographic semiorders of Tversky (1969), which violate transitivity, and the intersection quasi-orderings of Sen (1997), which violate completeness.

modest aim seems appropriate in some cases. For instance, with normative uncertainty, we are concerned primarily with which options are permissible, and less so or not at all with how to rank the various impermissible options. And in the cases of counterfactuals, the best-system analysis of lawhood, legal interpretation, and interpretivism about the mind, we care only about which antecedent worlds are overall *most* similar and about which deductive systems or interpretations are overall *best*. However, this strategy is insufficiently general, for whenever a multi-dimensional concept is gradable, we need a ranking of alternatives in order to make sense of its comparative form. Moreover, there are a host of impossibility results for FCCRs (see e.g., Sen, 2017, ch. A2*), and while we will not attempt to prove a Khmelnitskaya-style impossibility theorem for FCCRs operating on disparate scale types, we suspect that one may be lurking.

Third, and even more radically, we might reject Pareto principles like Weak Pareto. For instance, we might appeal to the notion of a golden mean and hold that some dimensions have a “sweet spot” such that increases along that dimension are good up until we reach that sweet spot, beyond which further increases are bad (Chang, 2002a). But we think it better to retain Weak Pareto and deal with apparent counterexamples by reinterpreting the underlying dimensions so that they are of the form “proximity to the ideal amount of X,” rather than simply “amount of X.” See Hedden and Muñoz (forthcoming) for further discussion.

Fourth, we might reject Khmelnitskaya’s assumption that dimensions within different subgroups are not comparable to one another.⁸ This response would be of a piece with a lesson commonly drawn from Arrow’s impossibility theorem, according to which Arrow’s result demonstrates the need for interpersonal comparisons of welfare (see Baccelli 2023 for a nuanced discussion of this issue). Formally, this response can be implemented either by denying that there are multiple subgroups with disparate scales, or by denying Informational Invariance across Subgroups on the grounds that meaningful information about interdimensional comparisons is not preserved by arbitrary combinations of different subgroups’ invariance transforms. But we find this strategy unappealing. First, interdimensional

⁸Bale (2008), for example, provides a semantics on which “mixed” comparisons like *Esme is more beautiful than Einstein is intelligent* are made possible by the construction of a *universal scale* which can in principle be accessed by all comparatives. On Bale’s view, in measurement-theoretic terms, all dimensions are measurable on an *absolute* scale. For alternative treatments of these sorts of comparisons, see van Rooij (2011); Dorr et al. (2023).

comparability strikes us as less plausible than interpersonal comparability of welfare. After all, in the latter, we are at least comparing the same *kind* of thing, namely welfare. By contrast, it seems nonsensical to compare how much a person knows with how much pleasure she experiences, or how courageous a person is with how honest she is, or how free a country's elections are with how well-protected its civil liberties are. Second, interdimensional comparisons are especially difficult to make sense of when the dimensions are measurable on different types of scales. This also highlights the way in which Khmelnitskaya's impossibility raises a novel challenge for multidimensional concepts, beyond the challenge already raised by Arrow's theorem. There are possible escape routes from Arrow's theorem, like the appeal to interdimensional comparability, which seem less promising as escape routes from Khmelnitskaya's. We should also note that the strategy we ultimately defend in section 7 works as a response to the challenges posed by the impossibility theorems of both Arrow and Khmelnitskaya. We will proceed, then, on the assumption that the setup required by Khmelnitskaya's theorem can arise—that there can be multiple subgroups, each with one of the required scale types, whose dimensions are not comparable with one another.

There is, no doubt, much more to be said about these four strategies. But we simply note our pessimism and move on to more promising ones.

5 Continuity and Lexical Priorities

Khmelnitskaya's impossibility theorem says that any aggregation function that satisfies her various other axioms yields a strongly dictatorial subgroup—i.e. one whose strict preferences *and* *indifferences* are respected by the social ordering. No other dimensions can come into play, not even as tie-breakers. This violates

Strong Pareto For any utility vectors u and v , if $u_i \geq v_i$ for every dimension $i \in N$ and $u_i > v_i$ for some $i \in N$, then u is *Fer* than v .

We find Strong Pareto highly compelling, almost as much as Pareto Indifference and Weak Pareto. It follows from the idea that dimensions are directional, such that one alternative ranking higher than another on a given dimension counts in favor of the former being more *F* than the latter overall. Then, when one alternative is *Fer*

than another on some dimension and equally F on all others, the former fact breaks the tie. See Hedden and Muñoz (forthcoming); Hedden (2023) for discussion.

Given Khmelnitskaya's result, we know that there is no aggregation function defined on an unrestricted domain which satisfies Strong Pareto, Pareto Indifference, Independence of Irrelevant Alternatives, Continuity, and Informational Invariance across Subgroups when there are multiple subgroups with disparate scale types (ordinal, interval, or ratio).

One way to avoid strong dictatorship and respect Strong Pareto is to drop Continuity. Continuity is motivated by the idea that the ordering of utility vectors shouldn't be hypersensitive to arbitrarily small changes in the values assigned by various dimensions. It should be possible to compensate for a tiny change along some dimension via sufficient changes along other dimensions. This rules out giving *lexical* priority to some dimensions.

Various forms of lexical prioritization, however, have been defended for some multidimensional concepts. In the case of normative uncertainty, the view known as "My Favorite Theory" involves a system of lexical priorities, whereby one option is more choiceworthy than another just in case, among the moral theories that are not indifferent between the two, the one you're most confident in ranks the one higher than the other (Gustafsson and Torpman, 2014). With scientific confirmation, we might say that fit with the data has lexical priority over other theoretical virtues, at least if the data are not noisy (McMullin, 1993). And with similarity and counterfactuals, Lewis (1979) can be read as endorsing a system of lexical priorities among different respects of similarity. On these views, one dimension (or subgroup thereof) is *weakly* dictatorial, in the sense that its strict preferences, but not necessarily indifferences, always determine the overall ordering.

Having said that, lexical priorities are highly implausible for many multidimensional concepts. To take just one example, consider the best-system analysis of lawhood. Lewis (1973, 73) notes that "Simplicity without strength can be had from pure logic, strength without simplicity from (the deductive closure of) an almanac," so that lexical priority for one over the other will limit the laws to mere logical truths or else have them include all contingent generalizations whatsoever.

Interestingly, however, there are various aggregation rules which satisfy all of Khmelnitskaya's axioms other than Continuity, and even Strong Pareto, but which do not assign lexical priority to any particular dimension or subgroup. For exam-

ple, suppose there are two subgroups, G_1 and G_2 . The dimensions in G_1 are ordinally measurable and fully comparable with each other; those in G_2 are interval-scale measurable and unit-comparable with each other. Consider an ordering of utility vectors that alternates from the leximin rule for G_1 to utilitarianism for G_2 and back to leximin for G_1 : u is *F*er than v iff

- The lowest utility in u_{G_1} (that is, G_1 's subvector of utilities) is greater than the lowest utility in v_{G_1} , or
- The lowest utilities in u_{G_1} and v_{G_1} are equal but the sum of utilities in u_{G_2} is greater than the sum in v_{G_2} , or
- The lowest utilities in u_{G_1} and v_{G_1} are equal and the sums of utilities in u_{G_2} and v_{G_2} are equal but the second-lowest utility in u_{G_1} is greater than the second-lowest lowest utility in v_{G_1} ,
- The two lowest utilities in u_{G_1} and v_{G_1} are equal and the sums of utilities in u_{G_2} and v_{G_2} are equal but the third-lowest utility in u_{G_1} is greater than the third-lowest lowest utility in v_{G_1} ,
- And so on.

If we think of each subgroup as having its own ordering of utility vectors—leximin for G_1 and utilitarianism for G_2 —neither subgroup is even weakly dictatorial: u_{G_1} can be leximin-preferred to v_{G_1} but u less *F* than v because the lowest utilities in u_{G_1} and v_{G_1} are equal but u_{G_2} is utilitarian-preferred to v_{G_2} . This ordering, however, satisfies all of Khmelnitskaya's axioms except for Continuity.

Although dropping Continuity does not force a particular dimension or subgroup to have lexical priority over all others, we do not ultimately find this to be a promising escape from Khmelnitskaya's impossibility. Discontinuities of the kind manifested in the example above—e.g., where the slightest decrease in the lowest utility assigned by G_1 outweighs arbitrarily large increases along other dimensions within and outside that subgroup—strike us as extreme, at least for many multidimensional concepts. Moreover, if we retain all of Khmelnitskaya's axioms other than Continuity, and strengthen Weak Pareto to Strong Pareto, we are still forced to prioritize certain subgroups in seemingly arbitrary and extreme ways.

Consider the set \bar{U} of all utility vectors in which all dimensions within each subgroup have the same utility—i.e., all u such that for every subgroup G , $u_i = u_j$ for all dimensions i and j in G . (Note that if all subgroups are singletons, as we think likely for many multidimensional concepts, then \bar{U} is the set of all utility

vectors.) And define the *orthant* of a vector u as the set of all v such that u_i and v_i have the same sign (including zero) for every dimension i . (For example, in two-dimensional Euclidean space, there are nine orthants in this sense: the four open quadrants, the four half-axes in between them, and one containing the origin.) Even without Continuity, Khmelnitskaya's remaining axioms plus Strong Pareto require that all vectors within any given orthant of \bar{U} compare to the origin in the same, strict way (excluding, of course, the orthant containing the origin itself): that is, if u and v are distinct vectors in the same orthant of \bar{U} , then either both are *Fer* than the origin or both are less *F* than the origin.⁹

This is already somewhat bizarre, but suppose we also restore a seemingly mild weakening of Continuity, by requiring the ordering of utility vectors to be continuous at the origin but possibly nowhere else—i.e., whenever a vector u is more (or less) *F* than the origin, there is some neighborhood Z_0 about the origin such that u is more (less) *F* than any v in Z_0 . It then follows, given Khmelnitskaya's other axioms and Strong Pareto, that there must be some particular subgroup G with the following property: for any vector $u \in \bar{U}$, if u_i is positive for every dimension $i \in G$, then u is *Fer* than the origin, and if u_i is negative for every $i \in G$, then u is less *F* than the origin.¹⁰ This, too, seems to us a bizarre constraint on multidimensional aggregation in general. And if this quasi-dictatorial subgroup G has a merely ordinal or interval scale type, then its influence will be even stronger, since the sign of its utilities is arbitrary.

An alternative weakening of Continuity requires the social welfare ordering to be continuous only *within* each orthant, so that whenever u is *Fer* than v and both belong the same orthant, there are neighborhoods Z_u and Z_v of u and v such that every u' in Z_u is *Fer* than every v' in Z_v . Given Khmelnitskaya's other axioms and Strong Pareto, this condition entails that there must be a linear hierarchy of

⁹*Proof.* For any vector $u \in \bar{U}$, let $O(u)$ denote u 's orthant in \bar{U} ; that is, $O(u) = \{v \in \bar{U} \mid \text{sgn}(v_i) = \text{sgn}(u_i) \text{ for all } i \in N\}$, where $\text{sgn}(0) = 0$. Let $\mathbf{0} = (0, \dots, 0)$. Take any distinct $u, v \in \bar{U}$ such that $O(u) = O(v)$. Then $\mathbf{0} \notin O(u)$. Informational Invariance across Subgroups for our scale types implies that $u \succ \mathbf{0}$ iff $v \succ \mathbf{0}$ (where \succ is the ordering of utility vectors), since for every subgroup G there is a positive real number k_G such that $v_i = k_G u_i$ for all $i \in G$; this is an admissible transformation for all seven scale types. Since \succ is complete, we have either $u \succ \mathbf{0}$, $u \prec \mathbf{0}$, or $u \sim \mathbf{0}$. Suppose for *reductio* that $u \sim \mathbf{0}$. There must be some $u' \in O(u)$ that is Pareto-superior to u (i.e., $u'_i \geq u_i$ for all $i \in N$ and $u'_j > u_j$ for some $j \in N$). Since $u' \succ u$ by Strong Pareto, $u \sim \mathbf{0}$ would imply $u' \succ \mathbf{0}$. But $u' \succ \mathbf{0}$ iff $u \succ \mathbf{0}$, since $u' \in O(u)$. Thus either $u \succ \mathbf{0}$ and $v \succ \mathbf{0}$, or $u \prec \mathbf{0}$ and $v \prec \mathbf{0}$. (Compare Nebel 2023c, Lemma 1.) \square

¹⁰This follows from the proof of Nebel (2023c, Lemma 2), up to the final paragraph.

orthants within \bar{U} —that is, all of the orthants in \bar{U} are strictly ranked against one another in such a way that any vector in a higher-ranked orthant is *Fer* than any in a lower-ranked one (see Naumova and Yanovskaya, 2001, Theorem 4.1). This sort of hierarchy strikes us as unattractive.

To sum up: while dropping Continuity does make Khmelnitskaya’s other axioms jointly satisfiable without making one dimension or subgroup even weakly dictatorial, the remaining axioms still have highly restrictive consequences when combined with Strong Pareto.¹¹ These consequences become even more extreme if Continuity is weakened in seemingly mild and reasonable ways, rather than entirely jettisoned. We do not claim that dropping Continuity is unpromising in all cases. But for a more plausible and generalizable escape from Khmelnitskaya’s impossibility, we must look elsewhere.

6 Dimensionalism

Khmelnitskaya’s impossibility theorem relies on the idea that the at-least-as-*F*-as ordering of alternatives depends only on the utilities assigned by the underlying dimensions of *F*ness, and not on any other features of the alternatives or the profiles in which the utilities are assigned. As noted in section 4, this is standardly called *Welfarism*, since in the traditional setting of welfare economics, the underlying dimensions concern individuals’ welfares. In the more general setting of multidimensional concepts, we call it *Dimensionalism*.

In Khmelnitskaya’s theorem, Dimensionalism follows from the conjunction of Unrestricted Domain, Pareto Indifference, and Independence of Irrelevant Alternatives. As we saw in section 4, the Welfarism (or, for us, Dimensionalism) Theorem—says that, given Unrestricted Domain, Dimensionalism is equivalent to the conjunction of Pareto Indifference and Independence of Irrelevant Alternatives.

A particularly interesting challenge to Dimensionalism arises in the setting of multidimensional concepts. In her discussion of value pluralism and the threat of an Arrovian impossibility result, Hurley (1985) argues that the conjunction of

¹¹Another surprising consequence of Khmelnitskaya’s other axioms plus Strong Pareto is that, even without any modicum of Continuity, they make it impossible to treat any subgroups symmetrically: there can’t be distinct subgroups G_1 and G_2 such that profiles U and V are assigned the same ordering of alternatives whenever $U_i = V_j$ and $U_j = V_i$ for every dimension $i \in G_1$ and $j \in G_2$. This follows from the proof of Nebel (2023b, Theorem 2).

Unrestricted Domain and Independence of Irrelevant Alternatives, in particular, conflicts with the supervenience of the evaluative on the descriptive. Here we will focus on her case against Unrestricted Domain.

Suppose that some alternative states of affairs are actually ranked thus-and-so along the various underlying dimensions of value (or, in her terminology, the various evaluative criteria): total welfare, equality, freedom, and so on. According to Hurley (1985, 511), “a given evaluative criterion cannot ‘change its mind’ about given alternatives; if it seems to, either the criterion has changed ... or the non-criterial description of the alternatives, hence the alternatives themselves, have changed.” That is, the identities of the alternatives fix what descriptive features they have, and by supervenience those descriptive features fix how they are ranked along the various underlying dimensions of value. Therefore, consideration of any other ranking of alternatives along these dimensions—as required by Unrestricted Domain—involves either some irrelevant counterevaluative (and hence counterpossible) supposition or else an illicit shift in which dimensions and alternatives we are considering.

In a similar vein, Morreau (2015) objects to Unrestricted Domain in the context of scientific confirmation and aggregating competing theoretical virtues. While it is a contingent matter how theories rank along the dimensions of fruitfulness and fit with the data, it is a necessary matter how they rank along the dimension of simplicity, for theories are abstract entities which have their degrees of simplicity essentially. In Morreau’s terminology, this means that simplicity is *rigid*. We should then reject Unrestricted Domain, for if one theory is in fact simpler than another, then it is necessarily simpler, and we should not demand that our aggregation function be defined for profiles which say otherwise.

While we concede that rejecting Unrestricted Domain may be plausible for some particular multidimensional concepts, we do not think it is a fully general solution to the problem raised by Khmelnitskaya’s impossibility theorem. First, even in cases where some underlying dimensions are supervenient or rigid, there may still be epistemic reasons for imposing Unrestricted Domain. Even if it is a necessary matter which theories are simpler than which, or which states of affairs are more equal than which, we may not *know* which are simpler or more equal. There may then be reason to demand that aggregation functions be defined for all epistemically possible profiles of dimensional utility functions, even if some such profiles

are metaphysically impossible. (This was Arrow’s own motivation for Unrestricted Domain—we cannot know in advance what the individuals’ preferences might be, and we want our aggregation function to work no matter what they turn out to be.) Of course, we might still know enough about how alternatives are ranked by different dimensions to rule out *some* profiles of dimensional utility functions. But we might motivate Unrestricted Domain as a modeling idealization. For there may be no natural, easy-to-work-with domain condition that exactly captures our knowledge about which profiles may be actual. By contrast, Unrestricted Domain is a natural domain condition which is theoretically fruitful and which at least doesn’t rule out any profiles that are compatible with our knowledge.

Second, while some multidimensional concepts may have dimensions which display the kind of supervenience or rigidity emphasized by Hurley and Morreau, others may not. Take *democratic*. It is implausible to say that a country could not vary at all with respect to the fairness of its elections or its turnout rates without being a different country altogether. Indeed, even in Hurley’s example of value pluralism, we find implausible her suggestion that alternatives could not differ in their descriptive properties without being different alternatives. This may hold if they are maximally specific possible worlds. But if they are more coarse-grained entities like public policies or incompletely-described states of affairs, there is no reason to think that they have all of their descriptive properties essentially.

Third, and most importantly, rejecting Unrestricted Domain may be a Pyrrhic victory. This is because we can generate analogues of Khmelnitskaya’s impossibility theorem with a significantly restricted domain.¹² For the sake of argument, let’s suppose with Hurley that all dimensions of F_{ness} supervene on certain properties that alternatives have essentially. Suppose also that we know exactly how each alternative fares with respect to each dimension. In the present framework, this knowledge doesn’t get us to a single utility profile because, for any given profile, there are infinitely many others that are informationally equivalent to it. So let’s suppose that the domain of the aggregation function is, for some profile U , the equivalence class $[U]$ of profiles that are informationally equivalent to U . This

¹²See also Nguyen (2019). Suppose we have a domain on which one dimension is rigid (so that its ranking of alternatives is held fixed) but which is otherwise unrestricted. Nguyen shows that given such a domain, along with the Arrovian axioms of Weak Pareto, Independence of Irrelevant Alternatives, and Non-Dictatorship, it follows that for any triple of alternatives, if the rigid dimension does not determine the overall ranking of any pair of them, then some other dimension determines the overall ranking of all three. This result strikes us as unpalatable.

is the “enriched single-profile” framework of Roberts (1980b), who provides a general recipe for deriving results in this framework from their multi-profile analogues (see also d’Aspremont 1985; Blackorby et al. 1990).

This recipe has two crucial ingredients. First, we must assume that there is sufficient diversity in how the alternatives fare along the various dimensions. In particular, we must assume that for any utility vectors u, v, w , there are alternatives x, y, z and a profile $U \in [U]$ such that $U(x) = u$, $U(y) = v$, and $U(z) = w$. This “richness” assumption allows the single-profile framework to mimic the force of Unrestricted Domain. However, even with richness, Pareto Indifference and Independence of Irrelevant Alternatives are not enough to secure Dimensionalism. Instead, it has to be imposed more or less directly, via this second ingredient:

Strong Neutrality For any profiles $U = (u_1, \dots, u_n)$ and $V = (v_1, \dots, v_n)$, if $u_i(x) = v_i(a)$ and $u_i(y) = v_i(b)$ for every dimension $i \in N$, then x is at least as F_U as y iff a is at least as F_V as b .

Strong Neutrality requires us to ignore non-utility characteristics of the alternatives, as well as other features of the profiles in which those utilities are assigned. It thus strengthens Independence of Irrelevant Alternatives, which is equivalent to the restriction of Strong Neutrality to cases where $x = a$ and $y = b$. Given the richness assumption, Strong Neutrality is equivalent to the existence of an ordering on the set of all utility vectors that determines the ordering of alternatives assigned to each profile in $[U]$ (d’Aspremont 1985; Blackorby et al. 1990). The rest of Khmelnitskaya’s impossibility theorem then goes through unscathed.

Both of these ingredients can be questioned. For example, some might maintain that there are necessary connections between the different dimensions of some concept, so that the utility assignments of the dimensions can’t vary independently of one another in the way required by richness. It would seem unreasonable to be confident, though, that this is true of *all* multidimensional concepts.¹³ We don’t know of any compelling argument against richness at the level of generality with which we are concerned. So we will simply suppose that richness is satisfied, in order to focus on Strong Neutrality.

¹³We concede, however, that it may be true of *some*. For instance, if moral value has (in some context) the underlying dimensions of each person’s welfare along with equality between them, then we cannot vary the utility assignment of each dimension independently, since changes to one person’s welfare (holding the others’ fixed) will necessarily change the degree of inequality.

Is Strong Neutrality defensible? We think so. At a minimum, it seems plausible as a default assumption for multidimensional concepts. Insofar as it is natural to talk about some concept being multidimensional with such-and-such underlying dimensions, it seems that those underlying dimensions should wholly determine the applicability of the concept. For instance, if *democratic* is multidimensional and its underlying dimensions are, say, turnout rate, protection of civil liberties, and fairness of elections, then it seems that whether one country is more democratic than another should depend only on how they compare with respect to turnout rate, protection of civil liberties, and fairness of elections.

In the traditional setting of welfare economics, Strong Neutrality is controversial because it makes social betterness depend only on individuals' welfares, to the exclusion of other morally important considerations such as rights (Sen, 1979). But insofar as this is a good objection to Strong Neutrality, this is precisely because individuals' welfares may not be the only underlying dimensions of social betterness; protection of rights may be another. (As Blackorby et al. 1990 and Weymark 2017 observe, this objection to Strong Neutrality is essentially an objection to Pareto Indifference, which is implausible in the setting of welfare aggregation precisely because welfare is, intuitively, not all that matters.) But once we have agreed on what the underlying dimensions of a concept are, Strong Neutrality seems highly plausible.

Hurley (1985, 517-8) objects to Strong Neutrality on the grounds that extra-dimensional information may affect the overall ranking of alternatives by affecting the relative weights of the underlying dimensions. In her toy theory, overall moral value depends on three underlying dimensions: total welfare, equality of resources, and equality of welfare. And the best state of affairs is the one that maximizes total welfare,

except in circumstances when by doing so we would leave someone with less than a certain minimum level of resources or would leave a handicapped person with less than a certain minimum level of welfare ... in these circumstances the theory tells us to distribute resources so as to maintain a certain roughly equal minimum level of resources for everyone and a certain roughly equal minimum level of welfare for the handicapped.

Here, the extra-dimensional information that someone is below the minimum level

(whether of resources or welfare) is relevant because it tells us that in these circumstances, the dimensions of equality get more weight than they otherwise would. This sort of extra-dimensional information can play a role reminiscent of the role allegedly played by so-called exclusionary, enabling, or modifying reasons, which affect the weights of various other reasons without themselves pointing in favor of certain alternatives over others (Raz, 1975; Dancy, 2004; Cullity, 2013).

Hurley (1985, 519) considers and rejects an obvious reply, namely that insofar as we are tempted by the toy theory she sketches, this is because we are thinking that the extra-dimensional information she finds potentially relevant—how many people fall below some minimum threshold of welfare—“must in effect be functioning as an additional criterion,” or dimension, of value. We find Hurley’s rejection of this reply to be overly hasty. She writes that “At least part of the sense in which falling below a certain level of welfare is a bad thing is already captured by considerations of welfare.” True, but the welfare-involving dimensions of value she already allowed for were *total* welfare and *equality* of welfare. Perhaps we need to add a further welfare-involving dimension having to do with the number of people falling below some critical level.¹⁴

Morreau (2014, 1265) gives another toy example of a violation of Strong Neutrality, this time for scientific confirmation and competing theoretical virtues:

Our evaluation of scientific theories lacks neutrality if the various criteria can go together differently depending on which theories we are choosing among. There needn’t be anything unscientific about this. There is no reason why the relative importance of fit and simplicity should be the same, say, in ecology as it is in physics.

If the theoretical virtues should be weighted differently depending on the domains of the theories in question, then the ranking of theories in terms of overall credibility will depend on extra-dimensional information, namely information about

¹⁴There is a general risk, however, that positing additional dimensions so as to deal with apparent violations of Strong Neutrality will result in necessary connections between dimensions and thereby render Unrestricted Domain (or the analogous richness assumption, in the single-profile framework) implausible. Compare Hedden and Muñoz (forthcoming), who defend Strong Pareto against purported counterexamples by positing an additional dimension which supervenes on the others. It is important to bear in mind, though, that Unrestricted Domain is sufficient but not necessary for securing Dimensionalism, and hence for generating a Khmelnitskaya-style impossibility. See Weymark (1998) for discussion of Welfarism (or Dimensionalism) theorems using restricted domains.

which domains they concern. But that information cannot be reinterpreted as constituting an additional dimension of scientific confirmation—being a physical theory rather than an ecological one doesn't by itself make for greater confirmation.

While the possibility adverted to by Morreau would require a rejection of Strong Neutrality, we deny that this possibility is actual. Insofar as it appears that the different theoretical virtues receive different weights depending on the scientific fields of the relevant theories, this just indicates that we haven't yet uncovered the fundamental theoretical virtues that guide abductive reasoning generally. If simplicity carries greater weight in physics than in ecology, we should seek an explanation of this fact. And we think that the explanation of this fact, whatever it may be, will advert to more fundamental theoretical virtues which apply equally in both fields. (In seeking out these fundamental theoretical virtues, we think it would be useful to consider comparisons between distinct "theories of everything." Because these theories must, by definition, encompass all fields, they must be compared in terms of domain-general theoretical virtues.)

We suspect that the foregoing considerations extend to other multidimensional concepts and potential violations of Strong Neutrality. In a slogan: apparent non-neutrality indicates non-fundamentality. That is, insofar as we seem to have a violation of Strong Neutrality, that indicates that we haven't uncovered (all of) the truly fundamental underlying dimensions of our multidimensional concepts. Of course, we have done nowhere near enough to show that this is always the case. But at a bare minimum, we think Strong Neutrality will hold for at least some, if not all, multidimensional concepts. Coupled with Roberts's Strong Neutrality-based recipe for getting single-profile impossibilities from multi-profile impossibilities, this means that a single-profile analogue of Khmelnitskaya's impossibility theorem will hold for at least some multidimensional concepts. We must look elsewhere, then, if we want a generally applicable escape from impossibility.

7 Dimensioned Quantities, Invariance, and Degree Functions

While each of the strategies we have considered so far has some plausibility for at least some multidimensional concepts, we have found all of them to be problematic, or at least insufficiently general to apply to all such concepts. We now want to suggest a strategy that, we think, offers a fully general escape route from

Khmelnitskaya's impossibility theorem.

The strategy, which comes from Morreau and Weymark (2016) and Nebel (2021, 2022, 2023a), is based on a simple defect of the orthodox framework we have been using so far. As we explain below, it is impossible within this framework to distinguish between two intuitively distinct kinds of “transformations” with respect to which we might want our aggregation function to be invariant. We propose an alternative framework in which these kinds of transformations can be distinguished, and we explain how Khmelnitskaya's conditions can be reformulated in this framework. Within this new framework, the apparently innocuous Informational Invariance across Subgroups looks suspicious and undermotivated.

Recall that Informational Invariance across Subgroups requires the same ordering to be assigned to any two profiles that are deemed informationally equivalent, by virtue of being related by a certain kind of transformation. It is motivated by the thought that the ordering delivered by our aggregation function shouldn't vary with merely representational changes in the units of measurement (e.g., pounds or kilograms) we use for each dimension. But strikingly, given Dimensionalism, Informational Invariance across Subgroups imposes a restriction on the ordering assigned to any *given* profile:

Intraprofile Invariance For any profile U and alternatives x, y, x', y' , if for every subgroup $G \in \Gamma$ there is some $\phi_G \in \Phi_G$ such that $U_G(x') = \phi_G(U_G(x))$ and $U_G(y') = \phi_G(U_G(y))$, then: x is at least as F_U as y iff x' is at least as F_U as y' .

But, on the face of it, it's not obvious why Intraprofile Invariance should hold. Assuming the invariance transform ϕ is nontrivial, it brings about a real difference between how x and x' fare along the various dimensions and how y and y' fare along those dimensions, not a merely representational change in the scale on which those alternatives are evaluated.

The problem is that, as Morreau and Weymark (2016) observe, the orthodox framework cannot distinguish between real changes in welfare—or, in our case, real changes along the underlying dimensions—and merely representational changes in measurement scales.¹⁵ Here is a toy example. Suppose that there are three dimensions, corresponding to the length, width, and height of various objects. And

¹⁵Following Sen (1977, 1542), the invariance conditions are “unable to distinguish between (i) everyone having *more* welfare (better off) in some real sense and (ii) a reduction in the *unit* of measurement of personal welfares.”

suppose that some observers measure these objects and write down the numbers which they take to represent their measurements. Each observer's measurement corresponds to a profile. Now suppose that two profiles U and V are related by a uniform doubling of these numbers. There are two possible interpretations of this. The first is that one of the observers—who generates profile U , say—takes all of the objects to be twice as long, wide, and tall as the other observer—who generates profile V —does. The other interpretation is that the observers agree on all of the spatial properties of the objects but have used different scales—say, inches vs. half-inches. These are, intuitively, very different kinds of transformations. Though we might, in this example, want a *bigger than* ordering to be invariant to transformations of both kinds, it is far from obvious that all multidimensional orderings must be invariant to real changes along the underlying dimensions *merely* because they are invariant to changes in the unit of measurement.¹⁶

Morreau and Weymark observe that this is not just a problem for the invariance conditions. It applies to other inter-profile conditions, such as Independence of Irrelevant Alternatives. Suppose, for example, that U and V in our example above assign the very same length, width, and height numbers to two particular objects, x and y , while U doubles the numbers assigned to all other objects. This could either mean that the observers generating these two profiles agree on the spatial properties of x and y while disagreeing about those of all other objects (because they use the same scale), *or* that the observers agree on all of the other objects but disagree on x and y (because they use different scales). Independence of Irrelevant Alternatives requires the orderings assigned to these two profiles to compare x and y in the same way. But, again, even if this verdict seems correct for the particular exercise of ordering objects by how big they are, it seems highly suspicious to require this in general, when the profiles reflect real differences in how x and y fare along the relevant dimensions.

The problems identified in the previous two paragraphs suggest that Independence of Irrelevant Alternatives and Informational Invariance across Subgroups are only plausible on incompatible interpretations of the utility assignments contained in each profile. Independence of Irrelevant Alternatives seems plausible only when being assigned the same number along each dimension means being just as F along

¹⁶Nebel (2022, 31–2) gives an example involving exponential decay of a radioactive material's mass, where the invariance requirement seems to deliver the wrong result.

every dimension. Informational Invariance across Subgroups only seems plausible when being assigned the same number along each dimension *doesn't* mean being just as F along every dimension, because the scales have changed.

The basic problem lies in the nature of the profiles: they only tell us the *numbers* assigned by the dimensions to each alternative, not the units in which they are measured—or, more generally, what the numbers *mean*. To avoid this problem, Nebel (2021, 2022, 2023a) suggests an alternative framework that is formally just like the framework of aggregation functions but in which numerical utilities are replaced by the properties they are supposed to represent.¹⁷ When these properties are measurable on interval or ratio scales, we can think of them as *dimensioned quantities* like 5 kg or 2 °C—which are not to be confused with the numbers used to represent them by particular scales (e.g., 5 and 2). But, as Nebel (2023a) shows, this framework also accommodates merely ordinally measurable properties, which are not “quantitative” in any ordinary sense.

We can adapt Nebel’s framework to the present setting as follows. For each dimension i of F_{ness} , there is a set of possible *degrees* \mathbb{D}_i to which something can be F_i . These degrees are not numbers; otherwise they would automatically be comparable across dimensions.¹⁸ Each profile $D = (d_1(\cdot), \dots, d_n(\cdot))$ is an n -tuple of *degree functions*, one for each dimension, which assigns a possible degree along that dimension to each alternative—that is, $d_i : X \rightarrow \mathbb{D}_i$. An aggregation function, just as before, assigns an ordering to each profile in its domain. We assume an unrestricted domain and impose the Dimensionalism axioms of Pareto Indifference and Independence of Irrelevant Alternatives, which can be translated straightforwardly in terms of degrees rather than numerical utilities. These axioms ensure that alternatives are ranked according to a single ordering on the set of all *degree distributions*—i.e., n -tuples of degrees, one for each dimension.¹⁹

In this framework, the measurability and interdimensional comparability of dimensions within a subgroup are captured by specifying what we call a *subgroup*

¹⁷Our main points could also, we expect, be expressed in the “scale-dependent” framework of Morreau and Weymark (2016), but we will not attempt this translation here.

¹⁸Identifying degrees with numbers might be appropriate for dimensions measurable on an “absolute scale”—such as, perhaps, voter turnout understood as the percentage of eligible voters who actually vote. Following Khmelnitskaya, however, we are not considering absolute scales here.

¹⁹See Nebel (2023a) for a proof of the Dimensionalism/Welfarism Theorem in this framework. We use the term “distribution” rather than “vector,” because these objects do not necessarily live in a vector space.

dimensional structure. A subgroup dimensional structure is a set of degrees together with one or more relations on that set, which together satisfy axioms characteristic of the desired scale type. The key innovation here is that unlike Nebel (2023a), we consider cases where different subgroups of dimensions have different kinds of dimensional structures.

For simplicity, in the following discussion, we assume that all dimensions *within* a given subgroup are fully comparable to one another. We lose almost no generality in making this assumption because (with the exception of unit-comparable interval scales) for any subgroup whose dimensions are not fully comparable to one another, we lose no information by thinking of each dimension as its own very small subgroup with full interdimensional comparability. (We simply set unit-comparable interval scales aside in this section, but see Nebel 2023a, sec. 5.3, on “hybrid difference structures” for the general approach we would take towards those scale types.) Similarly, we assume with Khmelnitskaya that dimensions in *different* subgroups cannot be compared.

For example, suppose the dimensions in subgroup G are ordinally measurable (and, as just stated, fully comparable with one another). Then the subgroup dimensional structure for G consists of its set of degrees \mathbb{D}_G and a linear ordering \succeq_G of degrees from greater to less. (By “linear ordering,” we mean a transitive, complete, and *antisymmetric* binary relation: for any degrees $a, b \in \mathbb{D}_G$, $a \succeq_G b$ and $b \succeq_G a$ only if $a = b$.) For example, there might just be a set of degrees of beauty, ordered from greater to less. We call the pair $(\mathbb{D}_G, \succeq_G)$ an *ordinal structure* just in case it is order-isomorphic to the real numbers: that is, there is a bijection $f : \mathbb{D}_G \rightarrow \mathbb{R}$ such that, for all degrees $a, b \in \mathbb{D}_G$, $a \succeq_G b$ iff $f(a) \geq f(b)$. The conditions required for such an isomorphism are stated in Appendix B. It is doubtful that all intuitively “ordinal” dimensions of multidimensional concepts have this structure, but we make this assumption (and analogous assumptions for the other structures below) to keep the framework as close as possible to that of Khmelnitskaya, where Unrestricted Domain implies that every vector of real numbers is assigned by some profile to some alternative.

Suppose next that the dimensions in G are measurable on (fully comparable) interval scales. Then we will again have a set of degrees \mathbb{D}_G , but the ordering \succsim_G (which will not be antisymmetric) will be defined on the set of all *pairs* of degrees $(a, b) \in \mathbb{D}_G \times \mathbb{D}_G$. We write the pair (a, b) as ab . These pairs are *differences* between

degrees—i.e., how much greater one degree is than another—and the ordering tells us which differences are greater or smaller. For example, if we are considering degrees of desire, we can order pairs of degrees in terms of strength of preference—i.e., how much more one wants one thing rather than another. The structure $(\mathbb{D}_G \times \mathbb{D}_G, \succsim_G)$ must satisfy the axioms of a *difference structure* as laid out by Nebel (2023a, sec. 5.1), restated here in Appendix B. These axioms are necessary and sufficient for the degrees to be representable by all real numbers, with the ordering of pairs represented by the arithmetic differences between numbers. That is, there will be a bijection $f : \mathbb{D}_G \rightarrow \mathbb{R}$ such that, for all degrees $a, b, c, d \in \mathbb{D}_G$, $ab \succsim_G cd$ iff $f(a) - f(b) \geq f(c) - f(d)$. This function is unique up to positive affine transformation. In such a structure, the difference between degrees a and b can be classified as positive ($ab \succ_G aa$), negative ($ab \prec_G aa$), or neutral ($ab \sim_G aa$), represented by an arithmetic difference that is positive, negative, or zero. The ordering of degree differences also induces a corresponding ordering \succeq_G of degrees themselves: $a \succsim_G b$ iff $ac \succeq bc$ for all degrees c .

Suppose next that the dimensions in G are measurable on a *ratio* scale. Then we will have a *concatenation* operation \circ_G , which takes any two degrees in \mathbb{D}_G and returns a third. Intuitively, this concatenation operation “adds” degrees together. For example, with length, the concatenation operation takes any two lengths a and b (which we can think of as equivalence classes of line segments) and returns what is intuitively their sum (which we can think of as the equivalence class of line segments resulting from any segments with lengths a and b adjoined in one direction from endpoint-to-endpoint). Bykvist (2021) argues that all values have such concatenation operations, and Nebel (2023d) considers such operations for welfare in particular. A linear ordering \succeq_G will be defined on the set \mathbb{D}_G , which is assumed to be closed under concatenation. The structure $(\mathbb{D}_G, \circ_G, \succeq_G)$ must satisfy the axioms of an *extensive structure* as laid out by Nebel (2023a, sec. 4), which again are provided in Appendix B. These axioms are necessary and sufficient for the degrees to be representable by all real numbers on a scale that is unique up to similarity transformation, with concatenation represented by addition. That is, there is a bijection $f : \mathbb{D}_G \rightarrow \mathbb{R}$ such that, for all degrees $a, b \in \mathbb{D}_G$, $a \succeq_G b$ iff $f(a) \geq f(b)$, and $f(a \circ_G b) = f(a) + f(b)$. The representation preserves both the ordering of degrees and the concatenation operation. In such a structure, a degree d can itself be classified as positive ($d \circ_G d \succ_G d$), negative ($d \circ_G d \prec_G d$), or neutral ($d \circ_G d = d$); the

“sign” of each degree is respected by any additive representation of the structure.

Before moving any further, let us summarize the main differences between this “qualitative” framework and the orthodox framework of numerical utility functions. In the qualitative framework, numerical utilities are replaced by the degrees they represent. These degrees can be dimensioned quantities (like 5 kg) or properties of some other kind. Each subgroup of dimensions has a structure, which determines the kinds of comparisons that can be made in terms of those dimensions. For example, in order to say that a is twice as F_i as b is F_j , we need degrees of these dimensions to stand in ratios to one another, as is supplied by an extensive structure. (These ratios can be defined in terms of the concatenation operation: for example, the ratio of a to b is 2 iff $a = b \circ b$.) We cannot make such comparisons with a mere difference structure, but we can say, for example, that the *difference* in F_i -ness between a and b is twice the difference in F_j -ness between c and d . With a merely ordinal dimensional structure, neither degrees nor degree differences stand in meaningful ratios, but we can say that x is more or less F_i than y is F_j (assuming, of course, that i and j belong to a common subgroup).

Since we have assumed the degrees of all dimensions to map one-to-one onto the real numbers, there may seem to be no gain in this framework. But this framework does not give rise to the same ambiguity as the orthodox one: being assigned the same degree of F_i -ness always means being equally F_i , and being assigned a greater degree always means being more F_i . Furthermore, there is no need to impose an invariance condition on the aggregation function to characterize the type of scale on which each subgroup’s dimensions are measurable, since the relevant features of these scale types are fully specified by each subgroup’s dimensional structure. They wear their scale types on their sleeves—or, better yet, as tattoos.

All of Khmelnitskaya’s conditions can be translated into this qualitative framework so straightforwardly that we will not write them out here (see Nebel, 2023a), with two exceptions. First, our formulation of Continuity appealed to the idea of “neighborhoods” of utility vectors, which we defined in terms of “distances” between them. We have not, however, defined a metric of distance between degree distributions. And specifying an arbitrary metric would seem unnatural in a setting where some dimensions can have merely ordinal structure. Fortunately, we can define “neighborhoods” in a more general way while preserving the force of Continuity. For any subgroup G and distinct degrees $a, c \in \mathbb{D}_G$, the set of all de-

degrees $b \in \mathbb{D}_G$ lying strictly between a and c is an *open interval*. We can then define a neighborhood of u as any set containing all u' such that, for every dimension i , u'_i lies within some open interval containing u_i . Then Continuity, translated straightforwardly in terms of degree distributions, is well-defined, and its numerical analogue is equivalent to the version stated earlier.²⁰

Second, rather than specifying a distinct class of invariance transformations for each subgroup dimensional structure and imposing Informational Invariance across Subgroups, we can state a single principle with the same force. The key is that the invariance transformations for each scale type are *automorphisms* of their corresponding dimensional structures. An automorphism is a one-to-one mapping of a structure onto itself which preserves all of the relations in the structure. For example, if we take each degree in an ordinal structure and apply a strictly increasing transformation—that is, $\phi : \mathbb{D}_G \rightarrow \mathbb{D}_G$ such that $a \succ_G b$ iff $\phi(a) \succ_G \phi(b)$ for all $a, b \in \mathbb{D}_G$ —the ordering of transformed degrees is the same as the ordering of untransformed degrees. Similarly, if we take each pair of degrees in a difference structure $(\mathbb{D}_G \times \mathbb{D}_G, \succeq_G)$ and map it onto one that is k times bigger (the axioms of the structure ensure that this operation is well-defined), the structure is the same set of differences ordered in the same way. More generally, for each subgroup $G \in \Gamma$, G 's dimensional structure determines a set Φ_G of transformations $\phi_G : \mathbb{D}_G \rightarrow \mathbb{D}_G$ that are automorphisms of G 's structure. We can then define the *multidimensional* structure for our concept F as the product of the various subgroup structures, so that the automorphisms of the multidimensional structure are obtained by independently applying automorphisms of the subgroup structures. More precisely, where $\prod_{i \in N} \mathbb{D}_i$ is the set of all degree distributions, a transformation $\phi : \prod_{i \in N} \mathbb{D}_i \rightarrow \prod_{i \in N} \mathbb{D}_i$ is an automorphism of the multidimensional structure iff, for every subgroup G , there is an automorphism ϕ_G of G 's dimensional structure such that $\phi_i = \phi_G$ for every dimension $i \in G$. Thus, in the present framework, Informational Invariance across Subgroups is equivalent to

Automorphism Invariance If two degree profiles are related by an automorphism of the multidimensional structure, then they must be assigned the same at-least-as- F -as ordering.

Given (qualitative versions of) the Dimensionalism axioms, this is equivalent to the

²⁰This is because the Euclidean topology and the product topology on \mathbb{R}^n are equivalent (see, e.g., Munkres, 2000, 123).

analogous condition on the ordering of degree distributions: for any distributions u, v, u', v' , if there is an automorphism ϕ of the multidimensional structure such that $u' = \phi(u)$ and $v' = \phi(v)$, then u is at least as F as v iff u' is at least as F as v' .

Khmelnitskaya's impossibility theorem, in this qualitative setting, can be stated as follows: if there are at least two subgroups, which are either ordinal, difference, or extensive dimensional structures, then there is no aggregation function defined on an unrestricted domain which satisfies (the qualitative analogues of) Pareto Indifference, Independence of Irrelevant Alternatives, Weak Pareto, Continuity, Automorphism Invariance, and Weak Subgroup Nondictatorship. We can assure ourselves that the theorem is valid in this setting by representing the degrees with numbers and appealing to Khmelnitskaya's original theorem.

When Khmelnitskaya's axioms are translated into these qualitative terms, we believe that the invariance conditions—captured by Automorphism Invariance—are highly suspect. This is because in the present framework, a transformation from one degree profile to another is unambiguously not a merely representational change in scale, but rather a real change in the degree to which each alternative is F along one or more dimensions. This is true even of those transformations that preserve each subgroup's dimensional structure—that is, the automorphisms.

Now, it may be possible to advance various arguments in favor of Automorphism Invariance. For instance, one might give an epistemic argument to the effect that degree profiles related by an automorphism of the multidimensional structure are indistinguishable, and so Automorphism Invariance is needed in order for the correct ordering of alternatives to be knowable. Or one might give a metaphysical argument to the effect that degree profiles related by such automorphisms do not represent genuinely distinct possibilities after all. (This is analogous to certain claims made by “comparativists” about physical quantities; see Dasgupta 2013.) Or one might argue for Automorphism Invariance on the grounds that aggregation functions which satisfy it are simpler and more elegant than ones which don't.

There is much to be said about each of these arguments (see Nebel, 2023a, sec. 6). Here we simply note that these arguments for Automorphism Invariance all rely on claims which are far more contentious than the mere thought that aggregation functions should be invariant with respect to merely representational changes, such as that involved in the shift from pounds to kilograms.

Though we are happy to allow that Automorphism Invariance may be satis-

fied by some multidimensional concepts (e.g., *bigger than*), we see little reason to impose it as a fully general constraint on multidimensional aggregation—in particular, when different subgroups of dimensions have different structures (unlike the length, width, and height example). Without Automorphism Invariance, there is a wealth of possible aggregation functions that are compatible with the remaining axioms once Automorphism Invariance is rejected. Here is a toy example, just for illustration. Suppose that there is one ordinal structure (G_1), one difference structure (G_2), and one extensive structure (G_3). One possible ordering of degree distributions proceeds as follows. Choose a particular numerical representation $f : \mathbb{D}_{G_1} \rightarrow \mathbb{R}$, a positive degree difference $ab \in \mathbb{D}_{G_2} \times \mathbb{D}_{G_2}$, and a positive degree $c \in \mathbb{D}_{G_3}$. We can then rank distributions u and v by, for example, summing the numbers assigned to each of G_1 's degrees, the ratio of each of G_2 's degree differences between u and v to ab , and the ratio of each of G_3 's degrees to c . (The axioms of the relevant structures ensure that these ratios are well-defined.) The aggregation function induced by this ordering satisfies all of the axioms other than Automorphism Invariance. We do not claim that this is an *attractive* ordering for any multidimensional concepts. It is just an example to show how it is possible to avoid a strongly (or even weakly) dictatorial subgroup while satisfying the other axioms, once we drop Automorphism Invariance.

One might worry that dropping Automorphism Invariance leaves us with *too many* possible kinds of aggregation functions. In social choice theory, the invariance conditions have proven extremely useful for axiomatically characterizing interesting classes of aggregation functions. Without these conditions, the possibilities may seem too unconstrained for us to actually figure out how dimensions should be aggregated. But, in the present context, the wealth of possibilities is a virtue, not a vice. Given the diversity of multidimensional concepts of interest in philosophy and elsewhere, we should—at least, prior to substantive theorizing about any particular such concepts—want there to be many different eligible ways to aggregate these functions. To narrow down the range of possibilities, we will need to use new axioms beyond the invariance conditions, with justifications that may be specific to particular multidimensional concepts.

To clarify, the lesson of this section is not that multidimensional concepts cannot be modeled in the orthodox numerical framework of aggregation functions. But it does mean that, when using that framework, we must be careful to keep in mind

its limitations—in particular, its inability to distinguish between real and merely representational differences along dimensions. When we keep that limitation in mind, it seems to us that the key interprofile conditions used in Khmelnitskaya’s results (Independence of Irrelevant Alternatives and Informational Invariance across Subgroups given nontrivial classes of invariance transforms) can be reasonably rejected. We can then make sense of multidimensional concepts after all, simply by minding our units.

8 Conclusion

Multidimensional concepts are ubiquitous, and they figure centrally in key debates in ethics, epistemology, and elsewhere in philosophy, not to mention economics, political science, biology, and other fields. How can the underlying dimensions of a multidimensional concept F be aggregated to yield meaningful verdicts on which things are *Fer* than which *overall*? In our view, it is fruitful to address this question using tools from social choice theory.

Our focus has been on a specific threat to the possibility of aggregation: different underlying dimensions of a given multidimensional concept may be measurable on disparate scale types, and Khmelnitskaya’s impossibility theorem suggests that sensible aggregation across scale types may be impossible.

We have argued that the best response to this impossibility is to adopt a qualitative framework in which numerical utilities are replaced by the degrees they are supposed to represent. In this framework, the (analogue of the) invariance conditions crucial to Khmelnitskaya’s impossibility theorem are ill-motivated. Once we reject them, we open up space for a wide range of possible aggregation functions.

We do not defend any particular aggregation function as the correct one, whether for any particular multidimensional concept or for all of them. We think that different multidimensional concepts may work differently in this regard (though learnability considerations may prevent them from working *too* differently), and that it may even be indeterminate which aggregation function applies to a given concept. We leave exploration of these possibilities to future research.

A Proof of Khmelnitskaya’s Impossibility Theorem

Suppose the set of dimensions N can be partitioned into $m \geq 2$ subgroups $\Gamma = \{G_1, \dots, G_m\}$, each of which has one of the following seven scale types: ONC, OFC, INC, IUC, IFC, RNC, RFC. Assume Unrestricted Domain, Pareto Indifference, and Independence of Irrelevant Alternatives. By the standard Welfarism Theorem of social choice theory, there is a unique ordering \succsim on the set of all utility vectors—the n -fold Euclidean space \mathbb{R}^n —such that, for any profile U and alternatives x and y , x is at least as F_U as y iff $U(x) \succsim U(y)$ (Bossert and Weymark, 2004, Theorem 2.2). We show below that if \succsim satisfies Weak Pareto, Continuity, and Informational Invariance across Subgroups, then there must be a strongly dictatorial subgroup, in violation of Weak Subgroup Nondictatorship.

Proof. Assume that \succsim satisfies Weak Pareto, Continuity, and Informational Invariance across Subgroups. The proof has three steps.

Step 1. Let \bar{U} denote the subset of \mathbb{R}^n in which all dimensions within each subgroup have the same utility—that is, $\bar{U} := \{u \in \mathbb{R}^n \mid \text{for every } G \in \Gamma, u_i = u_j \text{ for every } i, j \in G\}$. There is a bijection $z : \mathbb{R}^m \rightarrow \bar{U}$ such that, for any m -vector $u^m \in \mathbb{R}^m$, subgroup $G \in \Gamma$, and dimension $i \in G$, $z(u^m)_i = u_G^m$. We can define an ordering \succsim^m on \mathbb{R}^m in the obvious way: for any $u^m, v^m \in \mathbb{R}^m$, $u^m \succsim^m v^m$ iff $z(u^m) \succsim z(v^m)$. Clearly, this ordering \succsim^m satisfies Continuity and Weak Pareto.

Consider any $u^m, v^m \in \mathbb{R}^m$ and any $k^m \in \mathbb{R}_{++}^m$. Let \odot denote the component-wise (or “Hadamard”) product of two vectors: that is, $(u \odot v)_i = (u_i)(v_i)$. We have $u^m \succsim^m v^m$ iff $z(u^m) \succsim z(v^m)$, and $u^m \odot k^m \succsim^m v^m \odot k^m$ iff $z(u^m) \odot z(k^m) \succsim z(v^m) \odot z(k^m)$. By Informational Invariance across Subgroups (for scale types ONC, OFC, INC, IUC, IFC, RNC, and RFC), we have $z(u^m) \succsim z(v^m)$ iff $z(u^m) \odot z(k^m) \succsim z(v^m) \odot z(k^m)$, so $u^m \succsim^m v^m$ iff $u^m \odot k^m \succsim^m v^m \odot k^m$. Thus, \succsim^m satisfies information invariance for an RNC scale.

According to Tsui and Weymark (1997, Theorem 6), an ordering on \mathbb{R}^m satisfies Continuity, Weak Pareto, and information invariance for an RNC scale iff it is strongly dictatorial—that is, there is a $\tilde{G} \in \Gamma$ such that for any $u^m, v^m \in \mathbb{R}^m$, $u^m \succsim^m v^m$ iff $u_{\tilde{G}}^m \geq v_{\tilde{G}}^m$ (see Nebel 2023c for a simpler proof). Suppose without loss of generality that $\tilde{G} = \{1, \dots, |\tilde{G}|\}$.

Step 2. Now consider the subset of \mathbb{R}^n that is constant for all dimensions outside of \tilde{G} . For any $G \subseteq N$ and $u \in \mathbb{R}^{|G|}$, let $\hat{u} = \max\{u_i \mid i \in G\}$ and $\check{u} = \min\{u_i \mid i \in G\}$. Given any $c \in \mathbb{R}$, let \mathbf{c}_G be the constant $|G|$ -vector of c 's.

Take any $u_{\tilde{G}} \in \mathbb{R}^{|\tilde{G}|}$ and $c \in \mathbb{R}_{++}$. Weak Pareto and Continuity together imply that $(\hat{\mathbf{u}}_{\tilde{G}}, \mathbf{c}_{N \setminus \tilde{G}}) \succ (u_{\tilde{G}}, \mathbf{c}_{N \setminus \tilde{G}}) \succ (\check{\mathbf{u}}_{\tilde{G}}, \mathbf{c}_{N \setminus \tilde{G}})$. Continuity then implies that, for some $\xi \in \mathbb{R}$, $(u_{\tilde{G}}, \mathbf{c}_{N \setminus \tilde{G}}) \sim (\xi_{\tilde{G}}, \mathbf{c}_{N \setminus \tilde{G}})$. Thus, for any $c' \in \mathbb{R}_{++}$, $(u_{\tilde{G}}, \mathbf{c}'_{N \setminus \tilde{G}}) \sim (\xi_{\tilde{G}}, \mathbf{c}'_{N \setminus \tilde{G}})$, by Informational Invariance across Subgroups (for scale types ONC, OFC, INC, IUC, IFC, RNC, and RFC), since for any such c' there is some $k \in \mathbb{R}_{++}$ such that $c' = kc$. Since c' can be arbitrarily small, Continuity implies that $(u_{\tilde{G}}, \mathbf{0}_{N \setminus \tilde{G}}) \sim (\xi_{\tilde{G}}, \mathbf{0}_{N \setminus \tilde{G}})$. By an exactly similar argument, there must be some $\xi' \in \mathbb{R}$ such that, for any $-c' \in \mathbb{R}_{--}$, $(u_{\tilde{G}}, -\mathbf{c}'_{N \setminus \tilde{G}}) \sim (\xi'_{\tilde{G}}, -\mathbf{c}'_{N \setminus \tilde{G}})$, and thus $(u_{\tilde{G}}, \mathbf{0}_{N \setminus \tilde{G}}) \sim (\xi'_{\tilde{G}}, \mathbf{0}_{N \setminus \tilde{G}})$ by Continuity. Thus, $(\xi_{\tilde{G}}, \mathbf{0}_{N \setminus \tilde{G}}) \sim (\xi'_{\tilde{G}}, \mathbf{0}_{N \setminus \tilde{G}})$ by the transitivity of \sim . However, by the result of Step 1, this can only happen if $\xi = \xi'$. Thus, the value of ξ depends only on the components of $u_{\tilde{G}}$. Similarly, for any $v_{\tilde{G}} \in \mathbb{R}^{|\tilde{G}|}$, there is a ζ such that $(v_{\tilde{G}}, \mathbf{d}_{N \setminus \tilde{G}}) \sim (\zeta_{\tilde{G}}, \mathbf{d}_{N \setminus \tilde{G}})$ for any $d \in \mathbb{R}$.

By the result of Step 1, $(\xi_{\tilde{G}}, \mathbf{c}_{N \setminus \tilde{G}}) \sim (\xi_{\tilde{G}}, \mathbf{c}'_{N \setminus \tilde{G}})$ and $(\zeta_{\tilde{G}}, \mathbf{d}_{N \setminus \tilde{G}}) \sim (\zeta_{\tilde{G}}, \mathbf{d}'_{N \setminus \tilde{G}})$ for any $c, d, c', d' \in \mathbb{R}$. Thus, $(u_{\tilde{G}}, \mathbf{c}_{N \setminus \tilde{G}}) \sim (u_{\tilde{G}}, \mathbf{c}'_{N \setminus \tilde{G}})$ and $(v_{\tilde{G}}, \mathbf{d}_{N \setminus \tilde{G}}) \sim (v_{\tilde{G}}, \mathbf{d}'_{N \setminus \tilde{G}})$. So $(u_{\tilde{G}}, \mathbf{c}_{N \setminus \tilde{G}}) \succ (v_{\tilde{G}}, \mathbf{d}_{N \setminus \tilde{G}})$ iff $(u_{\tilde{G}}, \mathbf{c}'_{N \setminus \tilde{G}}) \succ (v_{\tilde{G}}, \mathbf{d}'_{N \setminus \tilde{G}})$. This means that the ordering depends only on the utilities of dimensions in \tilde{G} when all others have the same utility.

Step 3. To extend this conclusion to all of \mathbb{R}^n , take any $u_{\tilde{G}} \in \mathbb{R}^{|\tilde{G}|}$ and $v_{N \setminus \tilde{G}}, w_{N \setminus \tilde{G}} \in \mathbb{R}^{n-|\tilde{G}|}$. By Weak Pareto and Continuity, $(u_{\tilde{G}}, \hat{\mathbf{v}}_{N \setminus \tilde{G}}) \succ (u_{\tilde{G}}, v_{N \setminus \tilde{G}}) \succ (u_{\tilde{G}}, \check{\mathbf{v}}_{N \setminus \tilde{G}})$, and $(u_{\tilde{G}}, \hat{\mathbf{w}}_{N \setminus \tilde{G}}) \succ (u_{\tilde{G}}, w_{N \setminus \tilde{G}}) \succ (u_{\tilde{G}}, \check{\mathbf{w}}_{N \setminus \tilde{G}})$. By the result of the Step 2, however, $(u_{\tilde{G}}, \hat{\mathbf{v}}_{N \setminus \tilde{G}}) \sim (u_{\tilde{G}}, \check{\mathbf{v}}_{N \setminus \tilde{G}}) \sim (u_{\tilde{G}}, \hat{\mathbf{w}}_{N \setminus \tilde{G}}) \sim (u_{\tilde{G}}, \check{\mathbf{w}}_{N \setminus \tilde{G}})$, so $(u_{\tilde{G}}, v_{N \setminus \tilde{G}}) \sim (u_{\tilde{G}}, w_{N \setminus \tilde{G}})$. It follows that any n -vectors which have the same utilities for dimensions in \tilde{G} are in the same equivalence class under \sim , so the ordering on \mathbb{R}^n depends only on the utilities of those in \tilde{G} .

We can therefore define an ordering $\succ^{\tilde{G}}$ on $\mathbb{R}^{|\tilde{G}|}$ such that, for any $u_{\tilde{G}}, v_{\tilde{G}} \in \mathbb{R}^{|\tilde{G}|}$, $u_{\tilde{G}} \succ^{\tilde{G}} v_{\tilde{G}}$ iff, for any $w_{N \setminus \tilde{G}}, w'_{N \setminus \tilde{G}} \in \mathbb{R}^{n-|\tilde{G}|}$, $(u_{\tilde{G}}, w_{N \setminus \tilde{G}}) \succ (v_{\tilde{G}}, w'_{N \setminus \tilde{G}})$. This violates Weak Subgroup Nondictatorship.

□

B Dimensional Structures

This appendix lays out the conditions for ordinal, difference, and extensive structures as defined in section 7. The ordinal structure conditions are well-known. The other conditions, due to Nebel (2023a), are based on Hölder (1901); see Michell and Ernst (1996, 1997) for English translation and Michell (1999) for exposition. Nebel modifies Hölder's axioms with conditions from better-known structures defined by Krantz et al. (1971). Throughout this appendix, we drop the subgroup subscript G from the sets of degrees.

Ordinal Structures Let \mathbb{D} be a nonempty set of degrees and \succeq a linear ordering on \mathbb{D} . We call (\mathbb{D}, \succeq) an ordinal structure iff it is isomorphic to (\mathbb{R}, \geq) , for which the following conditions are individually necessary and jointly sufficient (see, e.g., Rosenstein, 1982, 37):

1. For every $b \in \mathbb{D}$, there are $a, c \in \mathbb{D}$ such that $a \succ b \succ c$.
2. There is a countable $B \subseteq \mathbb{D}$ such that for all $a, c \in \mathbb{D}$ with $a \succ c$, there is some $b \in B$ such that $a \succ b \succ c$.
3. For any partition of \mathbb{D} into subsets A and C such that $a \succ c$ for all $a \in A$ and $c \in C$, there is a $b \in \mathbb{D}$ such that $a \succeq b \succeq c$ for all $a \in A$ and $c \in C$.

Difference Structures Let \mathbb{D} be a set of degrees and \succsim an ordering on $\mathbb{D} \times \mathbb{D}$. We call $(\mathbb{D} \times \mathbb{D}, \succsim)$ a difference structure iff all of the following conditions are satisfied:

1. There exist $a, b \in \mathbb{D}$ such that $a \neq b$.
2. For all $a, b, c, d \in \mathbb{D}$, if $ab \succsim cd$, then $dc \succsim ba$.
3. For all $a, b, c, a', b', c' \in \mathbb{D}$, if $ab \succsim a'b'$ and $bc \succsim b'c'$, then $ac \succsim a'c'$.
4. For all $a, c \in \mathbb{D}$ there is a $b \in \mathbb{D}$ such that $ab \sim bc$.
5. For all $a, b, c \in \mathbb{D}$ such that $a \neq b$, there are unique $d, d' \in \mathbb{D}$ such that $cd \sim d'c \sim ab$.
6. For any partition of \mathbb{D} into subsets A and C such that $ac \succ ca$ for all $a \in A$ and $c \in C$, there is some $b \in \mathbb{D}$ such that $ad \succ bd \succ cd$ for all $a \in A, c \in C$, and $d \in \mathbb{D}$.

Extensive Structures Let \mathbb{D} be a set of degrees. Let \succeq be a linear ordering on \mathbb{D} . We assume that \mathbb{D} is closed under a binary operation $\circ : \mathbb{D} \times \mathbb{D} \rightarrow \mathbb{D}$. We call $(\mathbb{D}, \succeq, \circ)$ an extensive structure iff all of the following conditions are satisfied:

1. For all $a, b, c \in \mathbb{D}$, $a \circ (b \circ c) = (a \circ b) \circ c$.
2. For all $a, b, c \in \mathbb{D}$, $a \succeq b$ iff $a \circ c \succeq b \circ c$ iff $c \circ a \succeq c \circ b$.
3. There is some $a \in \mathbb{D}$ such that $a \circ a \succ a$.
4. For all $a \in \mathbb{D}$, if $a \circ a \succ a$, then there is a $b \in \mathbb{D}$ such that $a \succ b$ and $b \circ b \succ b$.
5. For all $a, b \in \mathbb{D}$, there are $c, d \in \mathbb{D}$ such that $a \circ c = b$ and $d \circ a = b$.
6. For any partition of \mathbb{D} into subsets A and C such that $a \succ c$ for all $a \in A, c \in C$, there must be a $b \in \mathbb{D}$ such that $a \succeq b \succeq c$ for all $a \in A, c \in C$.

References

- Accinelli, Elvio, and Leobardo Plata. 2008. Comparabilidad parcial con mediciones cardinales y elección colectiva. *El Trimestre Económico* 75:97–124.
- Adler, Matthew D. 2011. *Well-Being and Fair Distribution: Beyond Cost-Benefit Analysis*. Oxford University Press.
- Arrow, Kenneth. 1951. *Social Choice and Individual Values*. New York: Wiley.
- Bacelli, Jean. 2023. Interpersonal Comparisons of What? *Journal of Philosophy* 120:5–41.
- Bale, Alan Clinton. 2008. A Universal Scale of Comparison. *Linguistics and Philosophy* 31:1–55.
- Blackorby, Charles, David Donaldson, and John A. Weymark. 1990. A Welfarist Proof of Arrow's Theorem. *Recherches Économiques de Louvain / Louvain Economic Review* 56:259–286.
- Bossert, Walter, and John A. Weymark. 2004. Utility in Social Choice. In Salvador Barberà, Peter J. Hammond, and Christian Seidl, editors, *Handbook of Utility Theory: Volume 2 Extensions*, 1099–1177. Springer US.
- Broome, John. 2004. *Weighing Lives*. Oxford University Press.
- Bykvist, Krister. 2021. Taking Values Seriously. *Synthese* 199:6331–6356.
- Chang, Ruth. 2002a. *Making Comparisons Count*. Routledge.

- . 2002b. The Possibility of Parity. *Ethics* 112:659–688.
- Cullity, Garrett. 2013. The Context-Undermining of Practical Reasons. *Ethics* 124:8–34.
- d’Aspremont, Claude, and Louis Gevers. 1977. Equity and the Informational Basis of Collective Choice. *The Review of Economic Studies* 44:199–209.
- D’Ambrosio, Justin, and Brian Hedden. forthcoming. Multidimensional Adjectives. *Australasian Journal of Philosophy*.
- Dancy, Jonathan. 2004. *Ethics Without Principles*. Oxford University Press.
- Dasgupta, Shamik. 2013. Absolutism vs Comparativism About Quantity. *Oxford Studies in Metaphysics* 8:105–150.
- d’Aspremont, Claude. 1985. Axioms for Social Welfare Orderings. In *Social Goals and Social Organizations: Essays in Memory of Elisha Pazner*, 19–76. Cambridge University Press.
- Dorr, Cian, Jacob M. Nebel, and Jake Zuehl. 2023. The Case for Comparability. *Noûs* 57:414–453.
- Dworkin, Ronald. 1986. *Law’s Empire*. Belknap Press of Harvard University Press, 1st edition.
- Goodman, Jeremy, and Bernhard Salow. 2023. Epistemology Normalized. *Philosophical Review* 132:89–145.
- Goodman, Nelson. 1972. Seven Strictures on Similarity. In *Problems and Projects*. Bobbs-Merrill.
- Grinsell, Timothy. 2012. Avoiding Predicate Whiplash: Social Choice Theory and Linguistic Vagueness. *Proceedings of SALT* 22:424–440.
- . 2017. *Semantic Indecision*. Ph.D. thesis, University of Chicago, Chicago, IL.
- Gustafsson, Johan E., and Olle Torpman. 2014. In Defence of My Favourite Theory. *Pacific Philosophical Quarterly* 95:159–174.
- Hattiangadi, Anandi. 2020. Radical Interpretation and the Aggregation Problem. *Philosophy and Phenomenological Research* 101:283–303.
- Hedden, Brian. 2016. Does MITE Make Right?: On Decision-Making Under Normative Uncertainty. *Oxford Studies in Metaethics* 11:102–128.
- . 2023. Parity and Pareto. Unpublished manuscript, Australian National University.

- Hedden, Brian, and Daniel Muñoz. forthcoming. Dimensions of Value. *Noûs* .
- Hurley, Susan. 1985. Supervenience and the Possibility of Coherence. *Mind* 94:501–525.
- Hölder, Otto. 1901. *Die Axiome der Quantität und die Lehre vom Mass*. Teubner.
- Khmelnitskaya, Anna B., 1996. Social Choice Problems with Different Scales of Individual Welfares Measurement for Different Subgroups of Individuals. In Peter Kleinschmidt, Achim Bachem, Ulrich Derigs, Dietrich Fischer, Ulrike Leopold-Wildburger, and Rolf Möhring, editors, *Operations Research Proceedings 1995*, Operations Research Proceedings, 252–257. Springer.
- , 2002. Social Welfare Functions for Different Subgroup Utility Scales. In Andranik S. Tangian and Josef Gruber, editors, *Constructing and Applying Objective Functions*, Lecture Notes in Economics and Mathematical Systems, 515–530. Springer.
- . 2010. Existence of a Dictatorial Subgroup in Social Choice with Independent Subgroup Utility Scales, an Alternative Proof. In Adrian Van Deemen and Agnieszka Rusinowska, editors, *Collective Decision Making: Views from Social Choice and Game Theory*, Theory and Decision Library C, 111–123. Springer.
- Khmelnitskaya, Anna B., and John A. Weymark. 2000. Social Choice with Independent Subgroup Utility Scales. *Social Choice and Welfare* 17:739–748.
- Krantz, David H., R. Duncan Luce, Patrick Suppes, and Amos Tversky. 1971. *Foundations of Measurement: Vol. I*. Academic Press.
- Kroedel, Thomas, and Franz Huber. 2013. Counterfactual Dependence and Arrow. *Noûs* 47:453–466.
- Kuhn, Thomas S. 1977. Objectivity, Value Judgment, and Theory Choice. In *The Essential Tension: Selected Studies in Scientific Tradition and Change*, 320–39. University of Chicago Press.
- Lewis, David. 1973. *Counterfactuals*. Cambridge, MA, USA: Blackwell.
- . 1974. Radical Interpretation. *Synthese* 27:331–344.
- . 1979. Counterfactual Dependence and Time’s Arrow. *Noûs* 13:455–476.
- Lin, Eden. 2014. Pluralism about Well-Being. *Philosophical Perspectives* 28:127–154.
- List, Christian. 2002. Intradimensional Single-peakedness and the Multidimensional Arrow Problem. *Theory and Decision* 52:287–301.

- . 2004. Multidimensional Welfare Aggregation. *Public Choice* 119:119–142.
- Luce, R. Duncan, David H. Krantz, Patrick Suppes, and Amos Tversky. 2014. *Foundations of Measurement*, volume 3. Dover.
- MacAskill, William, Krister Bykvist, and Toby Ord. 2020. *Moral Uncertainty*. Oxford University Press.
- Maclaurin, James, and Kim Sterelny. 2008. *What Is Biodiversity?* University of Chicago Press.
- McMullin, Ernan. 1993. Rationality and Paradigm Change in Science. In Paul Horwich, editor, *World Changes. Thomas Kuhn and the Nature of Science*, 55–78. MIT Press.
- Michell, Joel. 1999. *Measurement in Psychology: A Critical History of a Methodological Concept*. Cambridge University Press.
- Michell, Joel, and Catherine Ernst. 1996. The Axioms of Quantity and the Theory of Measurement. *Journal of Mathematical Psychology* 40:235–252.
- . 1997. The Axioms of Quantity and the Theory of Measurement. *Journal of Mathematical Psychology* 41:345–356.
- Morreau, Michael. 2010. It Simply Does Not Add Up. *The Journal of Philosophy* 107:469–490.
- . 2014. Mr. Fit, Mr. Simplicity and Mr. Scope: From Social Choice to Theory Choice. *Erkenntnis* 79:1253–1268.
- . 2015. Theory Choice and Social Choice: Kuhn Vindicated. *Mind* 124:239–262.
- Morreau, Michael, and John A. Weymark. 2016. Measurement Scales and Welfarist Social Choice. *Journal of Mathematical Psychology* 75:127–136.
- Munkres, James. 2000. *Topology*. Pearson College Div, 2nd edition.
- Naumova, Natalia, and Elena Yanovskaya. 2001. Nash Social Welfare Orderings. *Mathematical Social Sciences* 42:203–231.
- Nebel, Jacob M. 2018. The Good, the Bad, and the Transitivity of *Better Than*. *Noûs* 52:874–899.
- . 2021. Utils and Shmutils. *Ethics* 131:571–599.
- . 2022. Aggregation Without Interpersonal Comparisons of Well-Being. *Philosophy and Phenomenological Research* 105:18–41.

- . 2023a. Ethics Without Numbers. *Philosophy and Phenomenological Research* forthcoming.
- , 2023b. Extensive Measurement in Social Choice. Unpublished manuscript, Princeton University. Available on PhilArchive: <https://philarchive.org/rec/NEBEMI>.
- . 2023c. Strong Dictatorship via Ratio-Scale Measurable Utilities: A Simpler Proof. *Economic Theory Bulletin* 11:101–106.
- . 2023d. The Sum of Well-Being. *Mind* doi: 10.1093/mind/fzad019.
- Nguyen, James. 2019. The Limitations of the Arrowian Consistency of Domains with a Fixed Preference. *Theory and Decision* 87:183–199.
- Okasha, Samir. 2011. Theory Choice and Social Choice: Kuhn versus Arrow. *Mind* 120:83–115.
- Rachels, Stuart. 1998. Counterexamples to the Transitivity of *Better Than*. *Australasian Journal of Philosophy* 76:71–83.
- Raz, Joseph. 1975. Reasons for Action, Decisions and Norms. *Mind* 84:481–499.
- . 1985. Value Incommensurability: Some Preliminaries. *Proceedings of the Aristotelian Society* 86:117–134.
- Roberts, Kevin. 1980a. Interpersonal Comparability and Social Choice Theory. *The Review of Economic Studies* 47:421–439.
- . 1980b. Social Choice Theory: The Single-profile and Multi-profile Approaches. *The Review of Economic Studies* 47:441–450.
- Rosenstein, Joseph G. 1982. *Linear Orderings*. Number 98 in Pure and Applied Mathematics. Academic Press.
- Schoenfield, Miriam. 2014. Permission to Believe: Why Permissivism is True and What It Tells Us About Irrelevant Influences on Belief. *Noûs* 48:193–218.
- Sen, Amartya. 1970a. *Collective Choice and Social Welfare*. Holden-Day San Francisco.
- . 1970b. Interpersonal Aggregation and Partial Comparability. *Econometrica* 38:393.
- . 1977. On Weights and Measures: Informational Constraints in Social Welfare Analysis. *Econometrica* 1539–1572.
- . 1979. Personal Utilities and Public Judgements: Or What’s Wrong With Welfare Economics. *The Economic Journal* 89:537–558.

- . 1997. *On Economic Inequality*. Oxford University Press, enlarged edition.
- . 2017. *Collective Choice and Social Welfare*. Harvard University Press, expanded edition.
- Sepielli, Andrew. 2009. What to Do When You Don't Know What to Do. *Oxford Studies in Metaethics* 4:5–28.
- Tarsney, Christian J. 2021. Vive La Différence? Structural Diversity as a Challenge for Metanormative Theories. *Ethics* 131:151–182.
- Temkin, Larry S. 1993. *Inequality*. Oxford University Press.
- . 2012. *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*. Oxford Ethics Series. Oxford University Press.
- Tsui, Kai-yuen, and John A. Weymark. 1997. Social Welfare Orderings for Ratio-Scale Measurable Utilities. *Economic Theory* 10:241–256.
- Tversky, Amos. 1969. Intransitivity of Preferences. *Psychological Review* 76:31–48.
- van Rooij, Robert. 2011. Measurement and Interadjective Comparisons. *Journal of Semantics* 28:335–358.
- Weymark, John A. 1998. Welfarism on Economic Domains. *Mathematical Social Sciences* 36:251–268.
- . 2016. Social Welfare Functions. In Matthew D. Adler and Marc Fleurbaey, editors, *The Oxford Handbook of Well-Being and Public Policy*. Oxford University Press.
- . 2017. Conundrums for Nonconsequentialists. *Social Choice and Welfare* 48:269–294.