Dimensions *as* Quotients

Andrew Y. Lee
University of Toronto, Philosophy

Abstract

A *dimension* is—roughly—a way of varying with respect to a kind. For example, hue is a dimension of color. This paper develops a systematic analysis of the dimension relation. The core idea is that for D to be a dimension of F is for the space of D-values to be a *quotient* of the space of F-values. A quotient is—roughly—a way of collapsing distinctions while preserving structure. I explain how this analysis yields intuitive verdicts, recovers a variety of structural principles connecting dimensions and their kinds, and connects the metaphysical concept of a dimension to the universal mathematical concept of a quotient. Afterwards, I identify a structural duality between the dimension relation and the parthood relation (akin to the duality between subsets and partitions), and I develop an important but underexplored distinction between dimensions and "parameters."

Introduction

A *dimension* is—roughly—a way of varying with respect to some kind.[1] As examples, consider the following statements:[2]

a. **HUE** is a dimension of **COLOR**.
b. **SPEED** is a dimension of **VELOCITY**.
c. **LONGITUDE** is a dimension of **GEOLOCATION**.
d. **TIME** is a dimension of **SPACETIME**.
e. **EXTRAVERSION** is a dimension of **PERSONALITY**.

The concept of dimension shows up in all sorts of contexts—ordinary discourse, scientific models, and nearly every area of philosophy. Within recent philosophical work, dimensions have played important roles in metaphysics, mind, ethics, aesthetics, language, and history.[3] If we were to rank structural concepts by how universal or flexible they are, then DIMENSION would rival canonical concepts such as LOCATION, MAGNITUDE, or PART.[4]

Yet surprisingly, there has been little philosophical work on dimensions. Although the concept is frequently used, it's rarely analyzed. This leaves a lacuna in the literature. What, exactly, do we mean when we say that some D is a dimension of some F?

If our use of 'dimension' were undisciplined and unsystematic, then that question might not even be worth investigating. If a definition of 'dimension' could be offloaded to another field—such as mathematics or physics—then there may be no need for philosophical analysis. But I'll argue that the subject-matter of dimensions is ripe for philosophical exploration and promising in its payoffs. The goal of this paper is to develop an analysis of dimensions that's intuitive, natural, formalizable, and fruitful.

My central idea is that for a property D to be a dimension of a property F is for the space of D-values to be a quotient of the space of F-values. A quotient—informally—is a way

[1] Other terms with nearby meanings include 'degrees of freedom', 'direction of movement', 'parameter', 'coordinate', 'variable', and 'aspect'.

[2] There are some important differences across these examples. But I've purposefully chosen a diverse sample that all fall under the target of my investigation. In some of my other work—especially on dimensionalization—I explore some of the relevant differences.

[3] As examples, see Funkhouser [2006] (determinables), Green [2020] (the perception-cognition border), Kulvicki [2020] (pictorial semantics), Hedden & Munoz [2024] (value), Mason [2025] (fittingness of emotions), Strawson [2025] (physical reality), White [2025] (aesthetic judgment), Maley [2026] (analog representation), Melamedoff [2026] (Spinoza's metaphysics), Jacobson [forth.] (valence), Kirby [forth]. (aesthetics), and D'Ambrosio & Stoljar [ms] (consciousness).

[4] I'll later return to the comparison between dimensions and parts—I'll argue that there's a "metaphysical duality" between the dimension relation and the parthood relation that mirrors the mathematical duality between partitions and subsets. The basic idea is that dimensions and parts are the two basic ways of generating something simple from something complex.

of collapsing a structure to generate a simpler version of that structure. I'll argue that this analysis satisfies a number of desiderata for a theory of dimensions, and I'll identify a deep connection between the dimension relation and one of the most general concepts in mathematics.

Here's the plan: **§1** clarifies the target sense of 'dimension', **§2** outlines some structural principles as desiderata for a theory of dimensions, **§3** explains how the concept of a partition is a useful tool for thinking about dimensions, **§4** develops the core theory of dimensions as quotients, **§5** discusses formal properties of the dimension relation, **§6** argues that there's an elegant "metaphysical duality" between the dimension relation and the parthood relation, and **§7** draws a philosophically significant structural distinction between dimensions and "parameters." The **APPENDIX** explains—in mostly informal terms—the mathematical relationship between quotients and partitions.

§1 The Target Concept

Let's start with a basic distinction (I'll henceforth use '*wrt*' as shorthand for 'with respect to'):

a *dimension* of F	a way of varying *wrt* F
the *dimensionality* of F	the number of ways of varying *wrt* F

This distinction in terminology maps to a distinction in questions:

ways-of-varying?	What is it for some D to be a dimension of some F?
***n*-dimensionality?**	What is it for some F to be *n*-dimensional?

This paper focuses principally on the ways-of-varying question. In other work, I examine the *n*-dimensionality question in detail. These questions, of course, interact with each other—and I'll discuss some of those interactions later in this paper. But the questions can also be investigated separately, and my answer to the ways-of-varying question won't depend on any particular answer to the *n*-dimensionality question.

When we talk about dimensions and dimensionality, we're typically talking about properties.[5] For example, hue (a property) is a dimension of color (a property), and color is three-dimensional (a property of a property). This makes the following line of reasoning tempting:

[5] I'll assume that kinds are determinable properties. But my arguments work even if we assign kinds (or dimensions) a different ontological status. For example, we also sometimes ascribe dimensions to concepts (COLOR), adjectives ('colored'), nouns ('color'), and spaces (color space).

(1) Dimensions are *properties*.
(2) Kinds are *properties*.
(3) Dimensionality is a *property of properties*.
Hence, (C) Dimensions are *properties of properties*.

But this inference, though tempting, is incorrect. Consider:

Observation: Hue is a dimension of color. But hue isn't a property of color. Instead, hue is a property of *colored things*. When we say that D is a dimension of F, we aren't committing ourselves to the claim that F itself has a D-value, but instead that the bearers of F-values are also bearers of D-values.[6]

For any property F, the **values** of F are the ways of being F. For two individuals to differ with respect to F is for those individuals to differ in their F-values. I'll assume the values of F are its maximal determinates. For example, the values for color are maximally determinate colors (red$_{34}$, green$_{17}$, etc.).[7]

The values of any property generate a **space**, meaning a set (whose elements are the values themselves) equipped with some structure (capturing relations such as similarity and magnitude). For example, the similarity relations between colors and the magnitude relations between masses determine the structures of color space and mass space. These spaces can have various kinds of mathematical structures, including topologies, metrics, algebraic operations, or orderings.[8] The figure below illustrates the geometry of the color disk (which captures hue and saturation but abstracts away from brightness):

[6] If a property F is itself a property of properties, then any of its dimensions D will likewise be a property of properties. But this is the exception, rather than the rule.
[7] The values of F should be distinguished from the realizers of F. For example, every color value is multiply realizable, in the sense that there are multiple distinct microphysical states of affairs that can realize that color value. If a surface is red$_{34}$, and we make a minor alteration to the atomic configuration of that surface, then the surface will (probably) still be red$_{34}$. But this doesn't mean that red$_{34}$ itself has multiple values. Those distinct realizers don't make a difference to the way the surface is with respect to *color* (even if they make a difference to the way the surface is with respect to microphysical configuration). Typically, even the maximal determinates of a property will be multiply realizable.
[8] What explains the structure of a property space? On realist views, the structure is grounded in the nature of the property itself. On anti-realist views, the structure is grounded in our representations of the property. I'll speak like a realist throughout the paper, but nothing essential hangs on this. It's uncontroversial that color, mass, and other properties are associated with spaces of values, even if it's controversial how to understand the nature of those spaces.

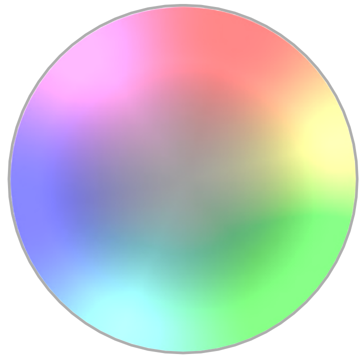

FIGURE 1: The color disk.

Throughout the paper, I'll sometimes move back and forth between using mathematical and metaphysical language. For example, I'll talk about partitions over properties, or distances between values, or one property being a quotient space of another. This is common in philosophical discussions of property spaces, and will streamline some of the prose.

Senses of 'Dimension'

Since the term 'dimension' is used in a number of different ways, it's useful to explain how my target sense of 'dimension' (ways of varying *wrt* some kind) relates to some other senses of 'dimension'.

Spaces: Alongside properties, it's common to take the bearers of dimensions to be spaces. As examples, the real line is one-dimensional, the Cartesian plane is two-dimensional, Newtonian space is three-dimensional, and Minkowski spacetime is four-dimensional. While spaces and properties are distinct ontological categories, there's a natural translation scheme when thinking about dimensions. For any space—whether physical, mathematical, or otherwise—there's an associated property: namely, location within that space. The dimensions of a space S will thereby correspond to the dimensions of the property location-within-S.[9] Conversely, whenever we talk of dimensions of a property, we can also talk about dimensions of the associated property space. For example, the property color is associated with the space of color values. Given this translation scheme, I'll freely move back and forth between talking about dimensions of F and dimensions of the space of F-values.

Dimensional Analysis: In physics, 'dimension' is most prominent in dimensional analysis, a method for analyzing the relationships between different physical quantities. In these contexts, the *dimension* of a quantity Q (*a*) specifies the basic quantities from which Q is derived, (*b*) abstracts over specific units of measurement, and (*c*) is expressed as a product

[9] A related sense of 'dimension' concerns the *spatial measurements* of an object: 'The dimensions of the box are 2"×12"×7.' Less often, 'dimension' is used to refer to "parallel planes of existence": 'This alien came from the 5th dimension!'. I think both these senses relate to the sense of 'dimension' I'm interested in, but I won't try to analyze the connections.

of powers. For example: 'The dimension of velocity is length over time (L/T).'[10] You might be tempted to think that the basic quantities in the dimensional analysis sense of 'dimension' just are the ways of varying in the variation sense of 'dimension'. But the example of velocity shows this isn't correct: neither length nor time are ways of varying with respect to velocity. Instead, the dimensions of velocity—in the sense I'm interested in—are speed and direction. I'll later argue that if D is a dimension of F, then anything with an F-value has a D-value. But notice that things with velocity values (namely material objects) aren't things that have length values (at least not in the relevant sense) or time values.

Dimensions vs. Parameters: In various philosophical contexts, including ethics and semantics, the expression 'dimension of F' has been used to mean (roughly) an input to the function that determines F-values. For example, if welfare is a function of pleasure, knowledge, and friendship, then pleasure, knowledge, and friendship have been called "dimensions of welfare." I'll later explain why this is distinct from my target sense of 'dimension', and I'll disambiguate by using the term 'parameter' to express this other notion. In §8, I'll discuss the relationship between dimensions and parameters in detail, and I'll explain how drawing the distinction clarifies a variety of philosophical questions.

Mathematics: In mathematics, the term 'dimension' has various definitions, depending on the kind of mathematical object under consideration.[11] But in each case, the definition specifies a measure over a class of mathematical objects, which outputs a number (specifying the dimensionality of that mathematical object). These definitions are obviously relevant to the *n*-dimensionality question.[12] But they have no obvious bearing on the ways-of-varying question. Even if we know how to measure the dimensionality of the space of F-values, we won't automatically know which properties are dimensions of F. To figure that out, we need a better understanding of the structural principles connecting dimensions and their kinds. That's the topic I'll turn to next.

§2 Ways of Varying

Occasionally, I'm asked what it is for something to be a dimension simpliciter.[13] But I find myself unsure of what is being asked. For the sense of 'dimension' I'm targeting, questions

[10] See Jalloh [2023] and Jacobs [2024] on dimensional analysis.

[11] The most prominent include definitions for *vector spaces* (number of basis vectors needed to generate the space), *topological spaces* (smallest n such that each point has a neighborhood homeomorphic to $\mathbb{R}^n$), *graphs* (smallest dimensionality of Euclidean space in which the graph, with unit edges, is representable), *partial orders* (smallest number of total orders whose intersection generates the partial order), and *rings* (longest chain of prime ideals).

[12] However, for reasons I discuss in other work (on the *n*-dimensionality question), even the answer to that question cannot simply be offloaded to mathematics.

[13] In some cases, the idea is that for D to be a dimension is for D to have a certain kind of intrinsic structure, such as magnitude or metric structure. But as I discuss later, there's arguably no kind of internal structure that's essential for D to be a dimension of F.

about dimensions should be relativized to kinds. Instead of asking the monadic question of what it is for something to be a dimension, we ought to instead ask the relational question of what it is for some D to be a dimension of some F.

I'll use **dimension relation** to express the relation in virtue of which one property is a dimension of another. To embark on our exploration of the dimension relation, let's start with some basic observations that can serve as desiderata for an analysis. From the fact that hue is a dimension of color, we might infer the following:

(1) Everything with a color value has a hue value.
(2) Nothing with a single color value has multiple hue values.
(3) All differences in hue are differences in color.
(4) Some colored things differ in hue values.
(5) Not all differences in color are differences in hue.

I believe these observations capture some of the structural connections between dimensions and kinds. These connections can be expressed schematically:

UNIVERSALITY	Everything with an F-value has a D-value.
INCOMPOSSIBILITY	Nothing with a single F-value has multiple D-values.
VARIABILITY	All D-differences are F-differences.
? ASPECTUALITY	Not all F-differences are D-differences.
? MULTIPLICITY	Some things with F-values differ in D-values.

I've put question marks beside MULTIPLICITY and ASPECTUALITY because those principles concern edge cases. In particular, MULTIPLICITY rules out dimensions that have just a single value, and ASPECTUALITY rules out dimensions that are structurally equivalent to their kinds. While both principles are intuitive, I think the choicepoint of whether to accept them ought to turn partly on questions about how to best formalize the dimension relation. As a comparison, consider how similar choicepoints arose for whether to define 'part' and 'subset' reflexively or irreflexively. I'll return to this point in §5.

Structural Principles

I'll now briefly motivate each principle mentioned above, continuing to use color as my main example.

UNIVERSALITY: Let *orange* be a determinate of color that corresponds to a proper subset of color space. Intuitively, orange isn't a dimension of color; instead, it's a determinate of color. If your initial intuitions diverge, then ensure that you aren't thinking instead of *orangeness* (meaning the degree to which an object is orange, which *is* intuitively a dimension of color) rather than orange (meaning a certain region of color space). Whereas every colored object has an orangeness value, not every colored object has an orange value. For example, something that's absolute black has an orangeness value (presumably zero), but it doesn't have an orange value (since absolute black isn't a value within the orange region of color space). Without UNIVERSALITY, we wouldn't be able to secure this result.[14]

INCOMPOSSIBILITY: This principle ensures that the values of a dimension (*a*) belong to the same category, and (*b*) don't overlap with each other. For example, INCOMPOSSIBILITY precludes red, sweet, and loud from each counting as distinct values along a dimension (since an individual can be red, sweet, and loud). It also precludes red_{34}, red, and colored from each counting as distinct values along a dimension (since red_{34}, red, and colored overlap). In fact, INCOMPOSSIBILITY arguably follows from a more general principle concerning properties: namely, that the maximally determinate values of a property must exclude each other.

VARIABILITY: Let hue × shape be the set of conjunctions of hue values and shape values. For example, the values of hue × shape include red ∧ square, red ∧ round, green ∧ square, and so forth. Intuitively, hue × shape isn't a dimension of color. This is because shape is independent of color. Hence, not all variations in hue × shape are variations with respect to color. If *x* is red ∧ round and *y* is red ∧ square, then *x* and *y* differ with respect to hue × shape, but *x* and *y* don't differ with respect to color. To satisfy the idea that dimensions are ways of varying with respect to a property, we need VARIABILITY.[15]

ASPECTUALITY and **MULTIPLICITY:** As mentioned earlier, ASPECTUALITY and MULTIPLICITY rule out edge cases. ASPECTUALITY captures the idea that dimensions are always more *coarse-grained* than their kinds. MULTIPLICITY captures the idea that dimensions must be ways of *varying*. In every paradigm example where some D is a dimension of some F, both ASPECTUALITY and MULTIPLICITY are satisfied. Yet as I noted earlier, there may nevertheless be some theoretical advantages to relinquishing these principles. For the time being, however, I'll continue treating these principles as desiderata for a theory of dimensions.

[14] You might worry that UNIVERSALITY precludes some ordinary ascriptions of dimensions. It's natural, for example, to say that cardiovascular strength is a dimension of health. Yet plants—which lack hearts—can still be healthy. However, I think such cases are better explained by quantifier domain restriction: the initial context is implicitly restricted to humans, rather than all organisms. If we were to give up UNIVERSALITY, then we would be faced with an unpalatable proliferation of dimensions. For example, just as plants don't have hearts, humans don't have roots, tentacles, or thoraxes, and so don't have values with respect to root structure, tentacle length, or thorax density. But it seems unreasonable to count all those as dimensions of health (at least if we're considering all organisms with health values).

[15] You might worry that VARIABILITY can't accommodate cases where changes along multiple dimensions cancel each other out. Consider, for example, an axiology where moral value is determined by welfare and equality, and where an increase in welfare can be exactly compensated by a decrease in equality, meaning that two worlds might be exactly the same in value yet have differing proportions of welfare and equality. This may appear to be a case where a D-difference (welfare) needn't yield an F-difference (value). But I'll later argue (in APPENDIX #2) that this conflates dimensions with what I'll call 'parameters', where the parameters of F are the inputs to the function that determines F-values. The axiology above is better classified as a view where value is unidimensional but has multiple parameters.

Intrinsic Structure

The principles above specified structural relations between dimensions and their kinds. But you might wonder whether there's any kind of intrinsic structure that's essential to dimensions. Suppose all we know is that D is a dimension of some F. What can we thereby infer about the structure of D?

Many invocations of dimensions implicitly assume some answer to this question. For example, it's frequently assumed that dimensions must be degreed, or have zero values, or have metric structure. But none of these kinds of structures is built into the concept DIMENSION. I'll illustrate using a few classes of counterexamples.

Multidimensional Dimensions: Some dimensions are themselves multidimensional. While this may initially sound slightly paradoxical, it aligns with our ordinary and theoretical uses of 'dimension'. For example, direction is a dimension of velocity, fitness is a dimension of health, and openness is a dimension of personality. Yet direction, fitness, and openness are each multidimensional. As a more philosophical example, consider two-dimensional semantics: it's obvious that both the primary and secondary dimensions of meaning are themselves multidimensional. Dimensions are ways of varying *wrt* a kind, and ways of varying can themselves be multidimensional.

Non-Linear Dimensions: Some dimensions have cyclical (rather than linear) structures. For example, hue—a dimension of color—is commonly taken to have the structure of a circle. And direction—a dimension of velocity—has the structure of a sphere. This means that there are no canonical orderings or privileged zero points amongst hue values and direction values. Therefore, dimensions needn't have degree structure or zero values.

Ordinal Dimensions: Some dimensions have ordinal (rather than interval, ratio, or absolute) structure.[16] For example, educational degree (BA, MA, PhD, etc.) might be taken to be a dimension of academic achievement. But there's no sense in asking whether the distance between PhD vs. MD is greater than the distance between MA vs. MSc. Educational degree has ordinal structure, rather than a stronger measurement scale.

Nominal Dimensions: Some dimensions have merely nominal structure, meaning the space of values is merely an unstructured set. For example, suppose we're deciding on which electronic files to transfer, and the relevant factors are file format (PDF, .docx, etc.) and file size (number of *kb*). Then it's natural to say that file format is one dimension of file type, even though there may be no additional structure on the space of file formats.

Categorical Dimensions: Some dimensions are categorical, meaning they have just two values (naturally thought of as "on"/"off"). For example, suppose we're US border agents deciding who to let into the country, and the only criteria guiding our decisions are (*a*) whether the person is a US citizen, and (*b*) the person's net worth. In such a context, it's

[16] See Stevens [1946] for the standard taxonomy on measurement scales.

natural to say that citizenship is a dimension of immigration qualifications, even though citizenship is a categorical property.

The General Lesson: If all we know is that D is a dimension of F, then we can infer almost nothing about the intrinsic structure of D-space. This is methodologically important. It indicates that the interesting structural properties of dimensions are extrinsic—relating dimensions and their kinds—rather than intrinsic. This also means that a general analysis of dimensions ought not assume that the properties have any particular kind of intrinsic structure. Instead, we need to think about abstract structural properties that are applicable to any kind of mathematical space.

§3 Partitions

Here's an observation that will serve as a stepping stone for my eventual analysis:

Observation: Whenever D is a dimension of F, D *partitions* the space of F-values. For example, hue partitions color: it's possible to divide color space into cells, where each cell consists of all color values associated with the same hue value.

A **partition** is a way of dividing a set into a collection of disjoint non-empty subsets such that each element is in exactly one of those subsets. Partitions are generated by an **equivalence relation** (a relation that's reflexive, transitive, and symmetric). Each equivalence class—called a "cell"—consists of exactly the elements that stand in that equivalence relation to each other. The diagram below depicts a partition on a disk-shaped space:

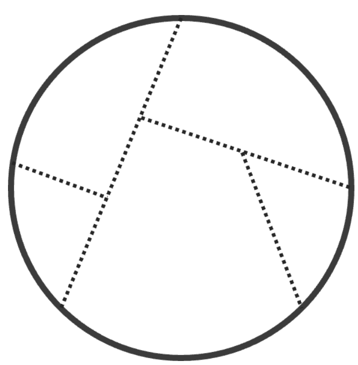

FIGURE 2: A partition.

As a mathematical example, the natural numbers $\mathbb{N}$ can be partitioned by *parity*, meaning whether a number is odd or even. Two natural numbers n and m stand in the relevant equivalence relation just in case either both are even or both are odd. Hence, this equivalence relation places every natural number into either the even cell or the odd cell. To denote the equivalence relation that generates a partition, it's standard to use the symbol ~. To denote the space that results from partitioning ℕ by ~, we write ℕ/~. Hence, ℕ/~ is the division of ℕ into two cells, "odd" and "even."

A property D partitions a property F just in case D generates an equivalence relation over the space of F-values. The relation is equivalence with respect to D. That is, F-values F_1 and F_2 are equivalent *wrt* D just in case F_1 and F_2 determine the same D-value.[17] For example, one way of partitioning color space is by equating all color values that are exactly the same in hue (abstracting away from differences in saturation and brightness). Each cell in the partition then corresponds exactly to a hue value. A partition over F, in general, is a way of abstracting over certain differences in F-values.

There's a close relationship between dimensions and partitions. Whenever D is a dimension of F, D generates a partition over F-space. This invites the following view:

The Partition Analysis
D is a dimension of F := D is a partition of F

The partition analysis is especially attractive when we see that it can satisfy *all* of the earlier criteria for a theory of the dimension relation. Consider the following relationships between the elements of a space (which correspond to F-values) and the cells of a partition (which correspond to D-values):

UNIVERSALITY	Each element is in exactly one cell.
INCOMPOSSIBILITY	No element is in multiple cells.
VARIABILITY	Different cells determine different elements.
MULTIPLICITY	Some elements are in distinct cells.
ASPECTUALITY	Not all elements are in distinct cells.[18]

Yet the partition analysis cannot be the end of the story. Partitions are defined in merely set-theoretic terms. But in nearly all cases of interest, both the kinds and the dimensions under consideration aren't merely sets, but instead **spaces**, meaning sets equipped with some structure. Moreover, whenever some D is a dimension of some F, the structure of D-space seems systematically related to the structure of F-space. A theory of dimensions

[17] I'll understand the entailment relation intensionally: for F-values to entail D-values is for it to be the case that necessarily, any two things with the same F-value must have the same D-value. This is equivalent to the claim that D-values metaphysically supervene on F-values. Accordingly, two F-values F_1 and F_2 entail the same D-value just in case necessarily, anything with F_1 has the same D-value as anything with F_2.

[18] To satisfy MULTIPLICITY and ASPECTUALITY, we need to exclude the trivial partitions. In particular, MULTIPLICITY requires excluding the trivially coarse partition that equates all F-values, while ASPECTUALITY requires excluding the trivially fine partition where each cell contains only a single F-value.

ought to explain how to think about these structural connections. This indicates that the partition analysis is still missing a key piece of the puzzle.

Nevertheless, the partition analysis is fundamentally on the right track. There's a natural way of strengthening it that satisfies the above desideratum. The solution lies in the concept of a quotient.

§4 Quotients

A quotient—informally—is a way of collapsing together elements to generate a simpler version of a space. Put another way, quotients abstract away from some distinctions while preserving a chosen kind of structure.

The term 'quotient' has a number of mathematical definitions, each local to a specific kind of structure. For example, there are distinctive definitions for groups, vector spaces, topological spaces, metric spaces, and partial orders. The mathematical details vary from case to case, depending on the kind of object being quotiented. But the common core is an operation that involves both abstraction and structure-preservation.

In a moment, I'll explain quotients in more detail. But let me first state my answer to the ways-of-varying question:

The Quotient Analysis

D is a dimension of F := D-space is a quotient of F-space.

I argued earlier that an analysis of dimensions ought to be applicable to a wide range of structures: discrete and continuous, linear and cyclical, geometric and algebraic. I've also argued that an analysis of dimensions ought to focus on the extrinsic structural relationship between dimensions and their kinds, rather than on the intrinsic structures of dimensions themselves. The quotient analysis—as we will see—satisfies both desiderata.

It's impossible, in this article, to explain the many mathematical details of quotients. But I'll focus on explaining enough for readers to understand the basic mathematical ideas behind the quotient analysis.

A Primer on Quotients[19]

Let's start with some core concepts relevant to any quotient construction:

[19] Interestingly, there are two approaches to defining quotients. The first is through ***structure-inheriting partitions***. The operation involves partitioning F-space and then inducing structure on that partition from the original F-space. The second is through ***structure-preserving functions***. The operation involves mapping F-space to D-space via a surjective structure-preserving map q that captures all the structure of D-space. These turn out to be mathematically equivalent up to isomorphism. Because of this, mathematical discussions of quotients

the **source space**	the space of F-values
the **quotient space**	the space of D-values
the **quotient map**	the structure-preserving function $q : \text{F} \longrightarrow \text{D}$

A quotient space D may be thought of as a collapsed version of its source space F. The collapse is defined by the quotient map q, a surjective (though usually not injective) function from F-space to D-space. The quotient map q induces a partition over F-space, where the equivalence relation ~ is defined as follows: $\text{F}_1 \sim \text{F}_2$ iff $q(\text{F}_1) = q(\text{F}_2)$. The following diagram illustrates these relationships via a quotient operation from color space to hue space:

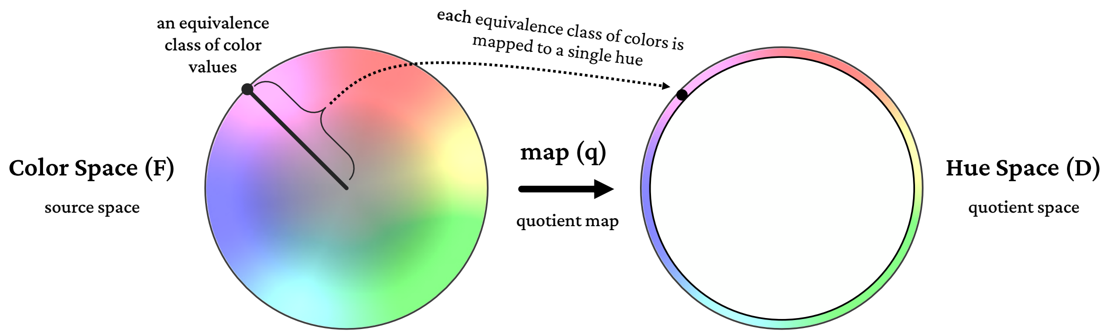

FIGURE 3: A quotient construction.

It's worth noting that the term **quotient** is polysemous. 'Quotient' can mean (**1**) the quotient space D, (**2**) the quotient map q, (**3**) the operation of applying q to F, resulting in D, or (**4**) the partitioned space F/~. These ambiguities are mostly harmless, since there are systematic relationships between these interpretations. When it's important to disambiguate, I'll use 'quotient space', 'quotient map', 'quotienting', or 'partitioned space'.

Quotients are closely related to partitions. Every quotient induces a partition. In fact, it's often useful to identify quotients by their corresponding partitions. In some contexts, quotients are even defined as partitions that inherit structure according to some principled rules.[20] To simplify the discussion, I'll focus here only on defining quotients in terms

sometimes often even treat these as two perspectives on the same mathematical object. To simplify the discussion, I focus only on the latter approach in this section. But the APPENDIX explains the former approach.

[20] There's a subtle metaphysical question about the relationship between the quotient space (D) and the partitioned space (F/~). Is each D-value literally an equivalence class of F-values? The answer turns on whether we draw distinctions between properties whose values are intensionally equivalent and whose spaces are isomorphic to each other. My official view is that dimensions are quotient spaces, but I'll leave open how these relate to partitioned spaces.

of structure-preserving maps (see the APPENDIX for a characterization of quotients via partitions).

Whether any arbitrary partition induces a quotient is a delicate question. The reason is that the concept of a quotient is structure-relative. When we ask whether D is a quotient of F, we need to specify the kind of structure under consideration in order to make the question precise. In the limit, if we care only about set-theoretic structure, then every partition yields a quotient (since for sets, quotients just are partitions). But for other kinds of structure—including algebraic operations, order relations, and geometries—arbitrary partitions may not preserve the relevant kind of structure.

This leads to a choicepoint for the quotient analysis. On a **restrictive version** of the view, only partitions that preserve the relevant kind of structure over F-space yield a dimension of F. On a **permissive version** of the view, every partition over F-space yields a dimension of F, though many of these partitions won't be structure-preserving.[21] The restrictive view makes the concept of a dimension more structurally interesting. But it also raises questions about exactly *which* kinds of structure must be preserved. It's not obvious how to answer this, since dimensions sometimes abstract away from one *kind* of structure while preserving another kind. For example, velocity has the structure of a vector space, and direction is a dimension of velocity, but direction lacks the structure of a vector space. Those skeptical that there's a principled way to define which kind of structure-preservation is needed may find themselves more drawn to the permissive view.

To understand quotients, we need to understand the concept of a structure-preserving map. But the definition of a structure-preserving map varies across mathematical domains. Common definitions include homomorphisms (algebraic structures), monotone maps (ordered sets), continuous maps (topological spaces), non-expansive maps (metric spaces), and linear maps (vector spaces). Because of this variance, the definition of 'quotient' also varies across mathematical domains. Nevertheless, these definitions are all unified by a universal definition of 'quotient', found in category theory.[22] The basic idea is that a quotient is a minimal way of collapsing a pair of structure-preserving maps.

A **quotient map** is a special kind of structure-preserving map. The intuitive idea is that a quotient map $q : \text{F} \to \text{D}$ must not only be structure-preserving, but must also capture all of the codomain's elements (meaning the function is *surjective*) and relations (meaning

[21] As far as I can tell, uses of 'quotient' in mathematics underdetermine these two senses. This is because nearly all discussions of quotients occur in specific mathematical contexts, where what's of interest is whether the structure associated with that mathematical category is preserved. Even in category theory, whether an object is a coequalizer may depend on the category under which one interprets that object.

[22] The relevant concept is that of a 'coequalizer', which generalizes the definitions of 'quotient' in specific mathematical categories (sets, vector spaces, topological spaces, rings, etc.). More technically, a *coequalizer* is a morphism $q : \text{Y} \to \text{Q}$ that's universal with respect to the following property: for some pair of morphisms $f : \text{X} \to \text{Y}$ and $g : \text{X} \to \text{Y}$, $q \circ f = q \circ g$. The intuitive idea is that a coequalizer of morphisms f and g is the minimal way of collapsing together f and g. See Mac Lane [1998] for a classic text on category theory.

the function is *final*—there is no additional structure in the codomain beyond those captured by the domain).[23] Put another way, every D-value must be the image of an F-value under q, and D-space must have no additional structure beyond that which is preserved under q. A quotient space can then be defined as follows:

D-space is a **quotient** of F-space	$=_{def}$	there exists a surjective, final structure-preserving map $q : \mathrm{F} \longrightarrow \mathrm{D}$

This primer—though brief—is enough to cover the core mathematical ideas. But to capture the whole story, I need to also say a few things about how the mathematics connects to the metaphysics.

From the Mathematics to the Metaphysics

There's an elegant way to bridge the mathematics with the metaphysics. The key insight comes from thinking about **relations**: a concept that will be familiar to both mathematicians and metaphysicians.

Every property space is structured by relations, such as relations of similarity and magnitude.[24] For example, the geometry of color space comes from the similarity relations between color values. Furthermore, both saturation values and brightness values stand in magnitude relations, which generates the ordinal structure of those spaces. And color qualities may also stand in other relations, such as vivacity or precision.[25]

This observation allows us to think about any property space as a **relational space**, meaning a set of elements endowed with some relations. The elements are the values of the property, and the relations are the kinds of higher-order relations expressed above. For relational spaces, structure-preserving maps are **relational homomorphisms**, meaning functions where relations in the domain are mirrored by relations in the codomain.[26] More specifically, a map q is a relational homomorphism just in case whenever a relation holds between some values $\mathrm{F}_1, \ldots, \mathrm{F}_n$, the corresponding relation holds between the images of those F-values under q. When this condition is satisfied, the map q "preserves the relations" over F-space.

[23] To make this precise, we need to also assign each space a *signature*, an ordered tuple of the structure (relations, operations, functions, etc.) on the set. This allows us to define which structure in the domain corresponds to which structure in the codomain. For example, the signature enables each F-relation to be matched with a corresponding D-relation.

[24] The limit case concerns property spaces that have merely set-theoretic structure. But even in those cases, we can define the set of relations over that space as the empty set.

[25] See Lee [2021] on a formal framework for modeling precision in quality-spaces.

[26] More technically, $q : \mathrm{F} \longrightarrow \mathrm{D}$ is a ***homomorphism*** iff for any relations R (over F-space) and R'(over D-space), $\mathrm{R}(\mathrm{F}_1, \ldots, \mathrm{F}_n) \longrightarrow \mathrm{R}'(q(\mathrm{F}_1), \ldots, q(\mathrm{F}_n))$.

Relational spaces are one of the most general ways of thinking about mathematical spaces. In fact, any mathematical space can, in principle, be characterized as a relational space.[27] And in philosophical work on "structuralism," structure is often characterized in terms of relations. Because of this, it's reasonable to think that the mathematical category of relational spaces is the kind of mathematical structure most relevant for the metaphysics of property spaces.

When mathematicians think about relational spaces, however, they aren't thinking about relations in metaphysical terms. In non-mathematical contexts, relations are polyadic properties: for example, 'taller than', 'more similar to', and 'between' all express relations. By contrast, the relations that occur in mathematical contexts are purely mathematical in nature. As an analogy, think about how points in mathematical spaces needn't be interpreted as any particular kinds of objects. Instead, points are dummy objects, individuated purely by the structural roles that they play.

If we understand quotients in a purely mathematical way, however, then we run into a problem with the metaphysics. The problem is this: any properties F and G will, from a purely mathematical point of view, stand in exactly the same quotient relations. Suppose, for example, that both duration and length have the structure of $\mathbb{R}+$. Then the space of duration values is isomorphic to the space of length values. But we obviously wouldn't want these properties to stand in exactly the same dimension relations.

The solution is to think of the relations as interpreted, rather than as purely formal. In other words, the relevant relations are the metaphysical relations that structure F-space and D-space, rather than purely mathematical relations that characterize the corresponding purely mathematical spaces. This ensures that isomorphic properties won't stand in exactly the same dimension relations. Whereas mathematicians care only about distinctions up to isomorphism, metaphysicians care also about the natures of the target elements and relations. By appealing to metaphysical relations between properties—such as those associated with similarity, magnitude, and precision—we can interpret quotients as capturing metaphysical (rather than only mathematical) structure.

This move enables us the quotient analysis to be sensitive to the nature—and not merely the structures—of the target properties. This enables us to recover the right dimension relations between properties. Furthermore, this allows us to better capture the sense in

[27] Usually, the relations that characterize relational structures are understood as on/off entities: for any n-adic relation and any n-tuple of individuals, either those individuals instantiate the relation or not. This leads to puzzles about how to relationalize geometrical structures, where there's a multitude of distinct relations that are all structurally related. For example, in a metric space, there may be infinitely many distinct distances, each of which will be associated with a distinct relation. I suspect that the solution is to think of relations—like any other property—as associated with spaces of values. For example, just as color is a property with many values (specific colors), distance is a relation with many values (specific distances). This approach to relational structures makes them a more powerful tool for uniting various kinds of mathematical structures. For limits of space, however, I won't develop this idea in detail.

which quotients involve abstracting away from a structure in order to generate a simpler version of that very same structure.

Applying the Analysis

To apply these ideas to a concrete example, let's consider how the quotient analysis works in the case of hue and color. Let q be a map from color space (F) to hue space (D), where each color value is mapped to its corresponding hue value. Intuitively, this map preserves the structure of hue while abstracting away from saturation and brightness. But does this satisfy the definition of a quotient?

Since every hue value will be the image of an equivalence class of color values (namely, all those that entail that hue value), it's obvious that q is surjective. Furthermore, the relations that structure hue space will always be applicable to the corresponding color values. For example, hue values stand in similarity relations that are naturally modeled by angular distance, measured by how far apart two hue values are along the hue circle. And these angular distance relations apply just as well to color values. Furthermore, whenever color values F_1 and F_2 stand in some angular distance relation, their images D_1 and D_2 will likewise stand in those same relations. In more intuitive terms, the similarity relations that capture hue space are mirrored by similarity relations in color space. These observations illustrate how the map q will be surjective, final, and structure-preserving. From this, we can conclude that q is a quotient map from color space to hue space.

It's easy to see that other examples of dimension and their kinds will follow a similar pattern. For example, consider speed and velocity. Suppose we partition the space of velocity by speed, meaning we equate all velocity values that are the same in speed. Then the partitioned space preserves the magnitude structure of speed while abstracting away from the angular structure of direction. Likewise, the map q from velocity values to speed values preserves magnitude relations while abstracting away from directional relations. Therefore, speed is a quotient of velocity.

In general, the quotient analysis of dimensions says that for D to be a dimension of F is for there to be a final surjective structure-preserving map q from F-space to D-space. But the relevant notion of structure-preservation concerns not only purely mathematical relations that characterize the mathematical structure of the relevant spaces, but also the metaphysical relations (such as similarity and magnitude) that define these property spaces.

The Quotient Analysis

You might wonder, at this point, what a counterexample to the quotient analysis look like. A natural strategy is to find a case where some D is intuitively a dimension of some F, yet where the space of D-values *cannot* be recovered by quotienting the space of F-values. This would require D to be structurally richer than F, in some respect. But in every intuitive example where some D is intuitively a dimension of some F, the structure of D-space is more coarse-grained than the structure of F-space. If this pattern is invariant across every example of a dimension, then it ought to be captured by an analysis of the dimension relation. The quotient analysis captures the pattern.

You might also wonder whether there are examples where D-space is a quotient of F-space yet where D isn't obviously a dimension of F.[28] I'm open, in principle, to a theory of dimensions that imposes additional restrictions. But the challenge is to find constraints that (*i*) apply across a comparably wide variety of mathematical structures: algebraic and geometric, discrete and continuous, linear and cyclical, and (*ii*) analyze the dimension relation in terms of a natural structural kind. Unless there's a plausible restriction that retains the generality and naturalness that makes quotients attractive in the first place, I'm inclined to stand by the quotient analysis.

It's worth noting that the quotient analysis doesn't appeal to notions such as essence, grounding, or naturalness.[29] However, it also doesn't preclude an appeal to such notions. Some might wish to appeal to these concepts to develop more metaphysically robust versions of the quotient analysis. For example, one might think that the structural relationship captured by the quotient analysis must hold in virtue of the essences of D and F.

As with the partition analysis, it's easy to verify that each of the structural principles expressed earlier—UNIVERSALITY, INCOMPOSSIBILITY, VARIABILITY, MULTIPLICITY, and ASPECTUALITY—can be satisfied. The first three principles follow straightforwardly from the quotient analysis (in analogous ways to the partition analysis). The last two principles turn on whether we permit the trivially coarse and trivially fine quotients—a point I'll return to in the next section. In my mind, it's remarkable that a universal mathematical concept can account for each of these structural principles. After all, the principles were developed as intuitive criteria that served as desiderata for a theory of dimensions. The fact that each criterion can be recovered by the quotient analysis illustrates the elegance of the view.

The key virtue of the quotient analysis (over the partition analysis) is that it captures the ways in which the structures of dimensions are systematically connected to the structures of their kinds. The quotient analysis captures these connections using a universal mathematical concept.

---

[28] One class of cases is where F is one-dimensional yet permits non-trivial quotients. For example, consider electric charge, which is often taken to have the structure of $\mathbb{R}$. It's possible to quotient $\mathbb{R}$ into three equivalence classes, which contain (1) all negative numbers, (2) zero, and (3) all positive numbers. The result is a quotient that inherits $\mathbb{R}$'s order structure (since positive > zero > negative). But that quotient still strikes me as a reasonable example of a dimension: it's natural to call polarity (positive vs. zero vs. negative) a dimension of charge.

[29] I suspect there's no universal answer as to whether dimensions or their kinds are more fundamental. Given this, I think an analysis of dimensions ought to let us to settle those questions on a case-by-case basis.

I began this paper by noting that dimensions are—roughly—ways of varying with respect to a kind. The quotient analysis allows us to take this rough idea and make it formally precise.

§5 The Dimension Relation

Suppose we accept the quotient analysis. We can now explore some of the formal properties of the dimension relation. To start, I want to consider two edge cases of dimensions, corresponding to the two trivial kinds of quotients. These can be defined by their corresponding partitions:

The Identity Quotient: $\forall F_i \forall F_j$: $F_i \sim F_j$ iff $F_i = F_j$
The Singleton Quotient: $\forall F_i \forall F_j$: $F_i \sim F_j$

Hence, the **identity quotient** (the trivially fine quotient) equates no distinct F-values, meaning that every F-value is its own equivalence class. Here the resultant quotient space is F-space itself. On the other hand, the **singleton quotient** (the trivially coarse quotient) equates every F-value, meaning there's only one equivalence class over the whole space. Here the resultant quotient space corresponds to a maximally determinate property that's instantiated just in case any F-value is instantiated. We can likewise define the corresponding "dimensions":

D is an **identity dimension** of F $=_{def}$ D-space is an identity quotient of F-space
D is a **singleton dimension** of F $=_{def}$ D-space is a singleton quotient of F-space

Should we include these in a theory of the dimension relation? The choicepoint corresponds to the question of whether to accept ASPECTUALITY and MULTIPLICITY. To accept ASPECTUALITY is to reject identity quotients (since those generate a quotient space that's equivalent to the original space).[30] To accept MULTIPLICITY is to reject universal quotients (since those generate a quotient space with just a single value).

To some extent, this is a verbal stipulation, rather than a substantive metaphysical stance. After all, no matter which we take to be the "official" definition, we can still draw the following distinctions:

[30] Actually, this turns on some further questions about property individuation. But I'll assume that if F and G are both intensionally equivalent and structurally isomorphic, then F = G. If we accept this, then proper dimensions can be defined as those that respect ASPECTUALITY and MULTIPLICITY, while improper dimensions include also dimensions that violate those principles.

D is a **proper dimension** of F	$=_{def}$	D is a non-trivial quotient of F
D is an **improper dimension** of F	$=_{def}$	D is a (possibly trivial) quotient of F

It may initially feel odd to think of the edge cases—involving identity quotients and singleton quotients—as counting as dimensions. But think about how it likewise feels odd when one first learns to think of the parthood and subset relations as reflexive. The motivations for the improper definitions of 'part' and 'subset' don't come from our ordinary judgments. Instead, the motivation is that those definitions yield more elegant formal frameworks. In my view, the question of how to best formalize the dimension relation ought to be treated analogously.

It turns out, in fact, that the improper definition yields a more elegant formal framework for dimensions. To start, the improper definition makes the dimension relation is reflexive, transitive, and anti-symmetric: the properties definitive of a partial order. **Reflexivity** holds because of the identity quotient. **Transitivity** holds because any composition of quotients is itself a quotient. **Anti-Symmetry** is more delicate: it requires a moderately coarse view of properties. But if we accept that properties that are both intensionally equivalent and structurally isomorphic are identical, then quotient symmetries will entail identity.

If we define the 'dimension' relation in this way, then every property F is associated with a **dimensional lattice**. A *lattice* is a partial order where every pair of elements has a unique supremum (least upper bound) and unique infimum (greatest lower bound). The elements of the lattice are dimensions of F. Equivalently, the elements can be thought of as quotients over F-space. The order structure is defined by the dimension relation: in particular, D ≤ F $=_{\text{def}}$ D is a dimension of F. The supremum is the identity quotient, and the infimum the singleton quotient. Here's a visual depiction of a very small lattice:

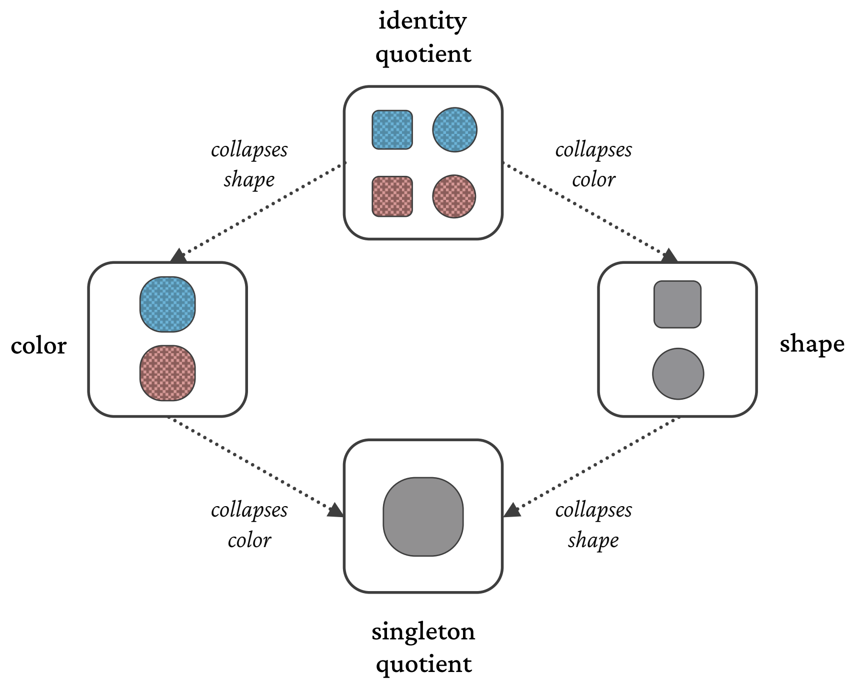

FIGURE 4: A dimensional lattice.

This formalization of the dimension relation associates every property with a hierarchy of dimensions, which come in varying levels of fineness and coarseness. However, this analysis might also invite a worry—which I turn to next—about the relationship between the quotient analysis of dimensions and the notion of dimensionality.

An Inequality

I previously distinguished the following two questions:

ways-of-varying? What is it for some D to be a dimension of some F?
***n*-dimensionality?** What is it for some F to be *n*-dimensional?

The quotient analysis is an answer to the ways-of-varying question. But you might worry that this answer has unpalatable implications for the *n*-dimensionality question. For example, color space is three-dimensional. But there are many ways to quotient color space. In general, the quotient analysis entails the following inequality:

INEQUALITY
dimensionality(F) ≠ number of D's such that D is a dimension of F

This may initially seem surprising. But in my view, it's a result we ought to embrace (even if we were to set aside the quotient analysis). A dimension, at first pass, is a way of varying with respect to a kind. But there are many cases where there are more than n ways of varying *wrt* an n-dimensional property. For example, it's hard to deny that both hue and redness are dimensions of color. But once we permit both of those, we ought to also permit saturation and brightness (since those accompany hue) and greenness, blueness, blackness, whiteness, and so forth (since those arguably stand or fall with redness). This means there are many properties—much more than three—that stand in the dimension relation to color. Yet it's uncontroversial that color is three-dimensional. The best resolution, in my view, is to accept INEQUALITY.

This doesn't mean that there's no connection between dimensionality and the dimension relation. The connection, in my view, can be clarified by thinking about **dimensionalization**, meaning the assignment of dimensions to a property. For example, hue, saturation, and brightness is one way of dimensionalizing color, while red/green, blue/yellow, black/white is another way of dimensionalizing color. While there are countless properties that are dimensions of color, any adequate dimensionalization of color will assign exactly three properties.

I explore questions about dimensionalization in more detail in other work. But here's a brief overview of the basic idea. To dimensionalize a property F is to take a step in constructing a model of F, meaning a mathematical representation of the space of F-values. The best models of F are those that optimize the balance between simplicity (roughly, number of dimensions postulated by the model) and strength (roughly, proportion of F-differences captured by the model). Since optimal models of color always postulate three dimensions, color is three-dimensional. But since there are many ways of quotienting color space, there are countless properties that stand in the dimension relation to color. Any of these properties may be selected when dimensionalizing color, though only collections of dimensions that maximize the ideal of orthogonality (defined in terms of degree of free recombinability between values) will be adequate dimensionalization of color.

§6 Dimensions vs. Parts

Here's an intriguing—though speculative—idea: the dimension relation is the "metaphysical dual" to the parthood relation. Both are "aspectuality" relations, meaning they concern relations between something complex and something simpler. But they capture two fundamentally different kinds of "aspects."

This speculative idea can be made more precise by thinking about the mathematical structures associated with each relation. I'll argue that this metaphysical duality—between dimensions and parts—is justified by a mathematical duality—between partitions (or quotients) and subsets (or subobjects). The picture that results generates a pleasing symmetry between these metaphysical concepts.

Partitions vs. Subsets

To illustrate, think about sets. There are two ways to "divide" a set. The first is by taking a **subset**. This cuts away a portion without collapsing any distinctions. The second is by taking a **partition**. This collapses some distinctions without cutting any portion.

Of course, there's a close relationship between these operations. Any subset canonically induces a partition: namely, a binary partition that divides a set into the subset and its complement. Any partition canonically induces a set of subsets, where each subset corresponds to a cell of the partition. But this merely illustrates the kind of mathematical duality I have in mind.

It turns out that this mathematical relationship runs much deeper. A partition, as I've explained, is closely related to a quotient. In fact, it's common to think of quotients as partitions that inherit structure from their source spaces. Therefore, dimensions are associated with partitions, while parts are associated with subsets. But quotients, as I've noted, are basically partitions with structure-inheritance. You might then wonder: is there a mathematical concept that stands to subsets as quotients stand to partitions?

Before I answer, let me raise another question. In mathematics, a concept that recurs across different contexts is **duality**. As examples, injective functions are dual to surjective functions, infimums are dual to supremums, the box operator is dual to the diamond operator, and conjunction is dual to disjunction. In category theory, the notion of duality can likewise be generalized, where this means (roughly) reversing the directions of the arrows in a definition. This definition of duality can recover the duality in local contexts, such as the ones expressed above. Now, I also mentioned earlier that the concept of a quotient has a universal definition, expressible in the language of category theory. This leads another question: is there a mathematical dual to the concept of a quotient?

The answer—to both questions—is a **subobject**. This means—roughly—an object that's contained within another object. This includes subsets (in set theory), subgroups (in group theory), subspaces (in topology), subposets (in order theory), and subgraphs (in graph theory). Just as quotients are basically partitions with structure-inheritance, subobjects are basically subsets with structure-inheritance.

The duality between quotients and subobjects can also be illustrated by thinking about the associated mappings. A *quotient map*, as noted previously, is a structure-preserving map that's both *surjective* (all elements in the codomain are mapped to) and *final* (there is no additional structure in the codomain). A related concept, when thinking about subobjects, is an *embedding*. This can be defined as a structure-preserving map that's both *injective* (distinct inputs map to distinct outputs) and *initial* (there is no additional structure in the domain). The following diagram depicts the difference between embeddings and quotients:

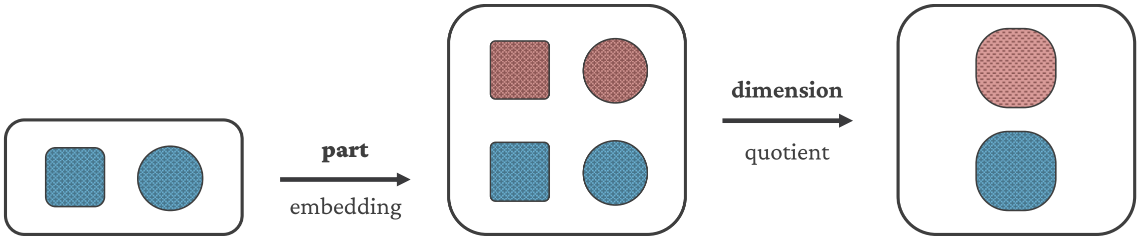

FIGURE 5: Embeddings vs. Quotients.

For limits of space, I won't go into the details of subobjects in this paper. But the analogies—as well as the explanation of how quotients relate to partitions—ought to be enough to see why this indicates a metaphysical duality between dimensions and parts.

Therefore, dimensions stand to partitions stand to quotients as parts stand to subsets stand to subobjects. This generates a pleasing symmetry between dimensions and parts, and an elegant way of connecting both of these basic metaphysical relations to general mathematical kinds.

Dimensions vs. Parts

The analogies between dimension and parts can be developed further. When thinking about *dimensions*, it's important to think also about *kinds* and *values*. What are the corresponding concepts for parthood? Here's a hypothesis:

dimensions	:	**kinds**	:	**values**
parts	:	**wholes**	:	**atoms**

The dimension relation relates dimensions and kinds; the parthood relation relates parts and wholes. Every property comes with a space of values (the basic units of properties that have no further determinates); every particular comes with a set of atoms (the basic units of particulars that have no further parts). For any kind F, we can define an associated domain of F-values; for any whole x, we can define an associated domain of x-atoms.

To make the comparison fully analogous, we would need to understand *atoms* in a relativized way, where atoms are relativized to wholes. Early on, I drew a distinction between values and realizers, where the values of a property can themselves be multiply realized. Similarly, instead of understanding an x-atom as an atom simpliciter, it might instead be understood as an atom relative to x. This mirrors the distinction expressed earlier between values and realizers, where the values of a property can themselves be multiply realized. Just as two objects can differ microphysically without differing *wrt* F, two objects can differ microphysically without differing *wrt* x.[31]

The table below depicts all the relationships I've discussed (dimension terms are in blue, parthood terms are in red, and neutral terms are in black):

	metaphysical relation		*unstructured mathematical relation*		*structured mathematical relation*	
aspect	**dimension**	**part**	**partition**	**subset**	**quotient**	**subobject**
complex	**kind**	**whole**	set		space	
unit	**value**	**atom**	element		element	

FIGURE 6: Dimensions *vs*. Parts.

These remarks express some initial comparisons between dimensions and parts. But there's much more to explore here. If my hypothesis is correct—if there are indeed these dualities within the metaphysics and the mathematics—then there's a whole research program to be developed on these connections between the mathematics and the metaphysics.

§7 Multiplicity Structures

Multidimensionality is sometimes conflated with other kinds of "multiplicity structures." Here's a brief list of definitions:

F is **multidimensional**	$=_{def}$	there are multiple ways of varying *wrt* F.
F-**pluralism**	$=_{def}$	there are multiple kinds denoted by 'F'.
F is **determinable**	$=_{def}$	there are multiple F-values.
F is a **genus**	$=_{def}$	there are multiple species of F.

These are all distinct.[32] In what follows, I'll briefly clarify the relationship between multidimensionality and the other kinds of structures. In the next section, I'll discuss a more subtle distinction between dimensions and "parameters."

[31] This isn't the standard way of understanding a mereological atom. But it strikes me as a sensible and underexplored way of thinking, and draws a closer analogy between dimensions and parts.

[32] But these structures are also sometimes conflated. For example, Hedden & Munoz [2024] seem to equate value pluralism, value multidimensionality, and value multiparametricity.

Pluralism

F-pluralism	$\nvDash$	F is multidimensional
F is multidimensional	$\nvDash$	F-pluralism

Just because there are multiple kinds of Fs doesn't mean there are multiple ways in which Fs can vary from each other. As an example, everyone agrees that velocity is multidimensional (speed and direction), but nobody thinks that there are multiple kinds of velocity. Hence, multidimensionality doesn't entail pluralism. As another example, almost nobody who endorses pluralism about truth thinks that truth is multidimensional (instead, even pluralists about truth tend to think that there are only two truth-values).

If F is multidimensional, then (by UNIVERSALITY) anything with an F-value has a value along each of the dimensions of F. But someone who endorses pluralism about F doesn't need to think that every instance of F is an instance of each kind of F. For example, one might be a value pluralist, in the sense that one thinks that multiple kinds of things (say, pleasure and knowledge) are valuable. But that doesn't mean that each instance of value has some degree of pleasure and some degree of knowledge.

Determinability

F is determinable	$\nvDash$	F is multidimensional
F is multidimensional	$\vDash$	F is determinable
F is determinate	$\vDash$	F is zero-dimensional

A *determinable* is a property that has multiple values, or *determinates*. For example, color is a determinable, and each specific color—red$_{34}$, green$_{17}$, etc.—is a determinate of color. If there are multiple ways of varying with respect to F, there must be at the very least two distinct values of F,[33] which entails that F is determinable. But F can be determinable without being multidimensional. For example, mass is determinable (there are many mass values) but unidimensional. Furthermore, questions about determinability underdetermine questions about multidimensionality. For example, the mere fact that color is a determinable (with each color value as a determinate) leaves open the dimensionality of color. In order to extract the dimensionality of color, we would need to know not only the determinates of color, but also the structural relations between those determinates.

[33] Why two values? Well, if F is multidimensional, then F must have at least two proper dimensions. Suppose, then, that D_1 and D_2 are both proper dimensions of F. Then D_1 and D_2 must each have at least two values. Now, if D_1 and D_2 each have two values, then it's natural to think that there must be at least four possible values of F (since $2 \times 2 = 4$). But since dimensions needn't be orthogonal, it's possible that both (1) the first value of D_1 is incompossible with the first value of D_2, and (2) the second value of D_2 is incompossible with the second value of D_2. That rules out two possibilities amongst the four permutations of D_1 and D_2 values, leaving two remaining possibilities. Hence, it's possible for F to be multidimensional while having only two values.

On the other hand, any property that's determinate has, by definition, only one value. But single-valued properties just are zero-dimensional properties. Hence, there's a close connection between being superdeterminates and zero-dimensionality. You might object by noting that some zero-dimensional properties—say, being a table—don't seem to be superdeterminate. But this objection conflates a property having a single value with a property having a single realizer. Nearly every property is multiply realizable, in that there are distinct states of affairs in virtue of which that property is instantiated. But not all those properties are multivalued: it's not the case that for any multiply realizable F, it's possible for *x* and *y* to vary *wrt* F.

You might wonder whether we could understand dimensions of F as (non-maximal) determinates of F. For example, redness is a dimension of color, and redness is also a determinate of color. But this trades upon two distinct interpretations of 'redness': (1) as a way of varying *wrt* color, vs. (2) as a region of color space. Consider, for comparison, hue (which doesn't admit of the same ambiguity). There's no sense in which hue is a determinate of color, since there's no subregion of color space that we could identify with hue. Furthermore, there's no subregion of color space that is devoid of hue values, since every color value is associated with hue values. Hence, while both dimensions and determinates are properties, it's never the case that a dimension of F can be identified with a determinate of F.

Genus/Species

S is a species of G	$\models$	$\neg$(S is a dimension of G)
D is a dimension of F	$\models$	$\neg$(D is a species of F)

In other words, the dimension relation is mutually exclusive from the 'species of' relation. To elicit some initial intuitions, consider how human is a species of (but not a dimension of) mammal, and how hue is a dimension of (but not a species of) color.

The reason is that distinct species of a genus exclude each other (human and walrus are both species of the genus mammal, but nothing is both a human and a walrus), while distinct dimensions of a property must be co-instantiable (since both hue and saturation are dimensions of color, anything that is colored must have *both* a hue value and a saturation value). More abstractly, if both B and C are species of A, then B and C must have disjoint extensions. By contrast, if both D_1 and D_2 are dimensions of F, then every individual with an F-value has both a D_1-value and a D_2-value.

There are also other differences between the genus/species relation and multidimensionality. Usually, both genuses and species are denoted using sortal terms, which can be combined with determiners ('the F', 'an F', etc.). For example, 'the mammal' or 'the walrus' is used to denote individual entities that fall under the category mammal or walrus. By contrast, we rarely use determiners to denote individuals that have values with respect to dimensions. For example, 'the hue' doesn't denote individual entities that have hue values, but instead hue values themselves.

§8 Dimensions vs. Parameters

I want to close by drawing an important but underexplored distinction between dimensions and "parameters." Let me start with an observation:

a puzzling observation

Many claim that ethical value is multidimensional because it's a function of welfare and equality.[34] Yet value, on such views, is often represented using scalars (the elements of a one-dimensional scale). The idea is that the dimensions of value (welfare and equality) aggregate to yield overall value (measured by scalars). But if value is best represented using scalars, then how could it also be multidimensional?

The puzzle is general. There are some FS where it's claimed both that **(1)** F is multidimensional, and that **(2)** F is scalar. But these claims seem straightforwardly contradictory: anything that has a scalar structure is, by definition, one-dimensional. To resolve the puzzle, we need to distinguish the concept of a dimension from the concept of a *parameter*.

Dimensions vs. Parameters

In mathematics, an argument of a function is sometimes called a **parameter**. A function with multiple parameters is thus a function that takes in multiple inputs. When thinking about mathematical functions, there's little risk of confusion between the number of parameters of a function vs. the dimensionality of its codomain. As examples, a multiparametric function—such as $f(x, y) = x + y$—might have $\mathbb{R}$ as its codomain, and a uniparametric function—such as $f(x) = (x, x)$—might have $\mathbb{R}^2$ as its codomain.

In philosophy, the term 'parameter' doesn't have a standard meaning. But it's worth coining the term, because there's a philosophical concept that plays an analogous role to the mathematical concept. Furthermore, a number of philosophical discussions use the term 'dimension' to express this concept.[35] That's the source of the puzzling observation from above. To dissolve the puzzlement, we should distinguish the following two notions:

D is a **dimension** of F	≈	D is a way of varying *wrt* F.
	:=	D-space is a quotient of F-space.

[34] See Hedden & Munoz [2024] as a prominent example.

[35] For some examples, see Hedden & Munoz [2024], D'Ambrosio & Hedden [2024], Hedden & Nebel [2024], and D'Ambrosio & Stoljar [ms].

P is a **parameter** of F	≈	P-values are inputs to the function that determines overall F-values.

The idea of a parameter is intuitive. Sometimes an individual's F-value is determined by its values along a collection of properties P_1, …, P_n.[36] For example, if moral value is determined by aggregating welfare and equality, then welfare and equality are each parameters of moral value. But from that, it doesn't follow that either welfare or equality are dimensions of moral value. In fact, if we grant that moral value is best represented using a one-dimensional scale, it's reasonable to infer that neither welfare nor equality are dimensions of moral value.

There are many examples that disentangle dimensions from parameters. Consider: (*a*) density is scalar, yet a function of mass and volume, (*b*) expected utility is scalar, yet a function of utility and probability, (*c*) land area is scalar, yet a function of shape and perimeter, and (*d*) GPA is scalar, yet a function of grades in each of one's courses. In each case, the target kind is typically represented using a subset of $\mathbb{R}$. That's evidence that the kind itself is one-dimensional, even though values along that kind are determined by multiple parameters.[37]

To evaluate whether a property is a parameter (and not a dimension) of F, it's useful to appeal to VARIABILITY: if D is a dimension of F, then D-differences are F-differences. Consider: two objects *x* and *y* can differ in both mass and in volume without differing in density. As long as the ratio of mass to volume is the same, *x* and *y* have the same density. Therefore, neither mass nor volume satisfy VARIABILITY *wrt* density. This indicates mass and volume are parameters, but not dimensions, of density. Analogous considerations apply to the other kinds mentioned above.

Readers who have previously used the term 'dimension' to mean parameter might wonder whether 'dimension' is polysemous, with one sense expressing parameters. I'm somewhat sympathetic to that thought: I think the fact that some academic

[36] I've purposefully left open how to understand the determination relation that relates parameters to their kinds. One option is to understand determination in modal terms. But another option is to appeal to a stronger relation, such as grounding or explanation. This strikes me as an attractive way to precisify the notion of a parameter, but I'll leave open which analysis is best. My main goal is not to develop a final analysis of parameters, but instead to explain how parameters are conceptually distinct from dimensions.

[37] What's the relationship between parameters and the sense of 'dimension' at stake in dimensional analysis? Well, the latter sense of 'dimension' is usually defined as a property of quantities, but parameters needn't be restricted to quantities. Also, the dimension of a quantity (in dimensional analysis) is expressed as a product of powers, but the parameters of F needn't be related to each other in such a way. Nevertheless, perhaps the notion of a parameter can generalize the sense of 'dimension' at stake in dimensional analysis.

subcommunities use ‘dimension’ to express parameters is enough to establish such a sense.[38] But the more relevant questions are which uses of ‘dimension’ are most prevalent and which ways of coordinating terminology are most fruitful. As far as I can tell, the most common use of ‘dimension’ expresses the notion identified at the beginning of this paper: a way of varying with respect to a kind. Many find that sentences such as ‘F is scalar / has the structure of $\mathbb{R}$ / varies in only one respect and F is also multidimensional’ initially jarring or confusing (at least outside theoretical contexts where parameters are salient). And in mathematics, ‘dimension’ is used to express measures of dimensionality—roughly, the number of ways of varying / degrees of freedom / directions of movement—rather than parametricity.

Nevertheless, my main aim is to draw a distinction and to illustrate its importance. The question of how to apply the label ‘dimension’ is verbal, but the distinction between dimensions and parameters is substantive. And for purposes of conceptual clarity, it’s desirable to use different terms for different concepts, especially when the concepts are closely related and easily conflated, and especially when sensitivity to the distinction can clarify some philosophical issues (as I’ll argue next). When ‘dimension’ is used in different senses that interact in complex ways, there is a risk of pervasive confusion, especially when the context underspecifies which sense is at stake. The considerations above are reason to think that the use of ‘dimension’ in the ways-of-varying sense is common, intuitive, and entrenched. That motivates searching for alternative terms for the nearby concepts. I’ve offered a solution: distinguish ‘dimension’ from ‘parameter’.

A Structural Analysis of Parameters

There are some interesting structural relationships between dimensions and parameters. By identifying these relationships, we can also clarify a number of philosophical issues.

To start, we need to make the notion of a parameter more precise. Just as with dimensions, we can formulate a number of structural principles that relate parameters to their kinds. Obviously, it’s beyond the scope of this paper to analyze parameters in as much detail as dimensions. But there are two principles that are useful to highlight here. These are analogues of structural principles on dimensions that we encountered earlier. Let P_F be a complete set of parameters for F.[39] Then:

[38] There are interesting questions about how to think about the meanings of theoretical terms when there is a distinctive pattern of use associated with a small body of researchers. See Lee & Mankowitz [2026] for discussion of this for the case of ‘conscious’ (in the phenomenal sense).

[39] What does ‘complete’ mean? Well, one answer is that a complete set includes *all* the parameters of F. But that might be overkill, since it might be the case that many parameters overlap (and so including all will be massively redundant). A more minimal option is to define a set of parameters P_F as complete relative to F just in case each P_F-value uniquely determines an F-value.

ASPECTUALITY*	Not all P_F-differences are F-differences.[40]
MULTIPLICITY*	Some things with F-values differ in P-values.

The idea behind MULTIPLICITY* is simply that parameters must have multiple values. The idea behind ASPECTUALITY* is that the function from parameters to F-values be *many-to-one*. More precisely, let $f: (P_1, \ldots, P_n) \rightarrow F$ be the function from the parameters of F to F-values. By many-to-one, I mean the conjunction of non-injectivity—multiple sets of P_i-values that yield the same F-value—and surjectivity—every F-value is the image of some set of P_i-values.

The invocation of ASPECTUALITY* and MULTIPLICITY* may remind you of the distinction between proper vs. improper dimensions. Given this, it's natural to draw an analogous distinction between proper vs. improper parameters. In what follows, I'll focus only on *proper parameters* and *proper dimensions* (omitting 'proper' in the prose).

These structural constraints make dimensions and parameters mutually exclusive. For parameters, the many-to-one function goes from the space of parameters to the space of F-values. For dimensions, the many-to-one function goes from the space of F-values to the space of D-values.[41] In this sense, dimensions stand in an inverse relationship to parameters.

Even if dimensions and parameters are mutually exclusive, some properties may be both multidimensional and multiparametric. Consider health: it's plausible both that there are multiple ways of varying *wrt* health and multiple parameters that determine something's health value. Perhaps one reason that multidimensionality and multiparametricity are sometimes conflated is because many properties fall under both categories. Other common examples, such as beauty, creativity, and athleticism, seem to exhibit similar patterns.

By distinguishing dimensions and parameters, we can sharpen our understanding of a variety of philosophical issues. To start, consider **aggregation**. Some philosophers have assumed that if F is multidimensional, then the dimensions of F must aggregate to yield overall F-scores.[42] However, many multidimensional properties aren't aggregative. For example, personality is multidimensional, but the dimensions of personality (such as openness and extraversion) don't aggregate to yield "overall personality scores." Even if you have

[40] Notice that ASPECTUALITY for parameters (not all P-differences are F-differences) inverts ASPECTUALITY for dimensions (not all F-differences are D-differences). By contrast, MULTIPLICITY has the same form for both dimensions and parameters. Also, the principles for parameters concern the whole parametric space (meaning the space of all parameters of F), while the principles for dimensions concern just a single dimension.

[41] For parameters, the many-to-one function is from the space of *all* parameters, meaning the space of all $(P_1, \ldots, P_n)$-values. By contrast, for dimensions, the many-to-one function is to the space of values for a single dimension. This illustrates that there are also some structural asymmetries between dimensions and parameters.

[42] See Hedden & Nebel [2024].

a higher value than me along every dimension of personality, that doesn't mean you have a "higher personality score" than me!

Instead, aggregation is a mark of multiparametricity. If F-values aggregate to overall F-scores, then there must be multiple parameters that are combined to determine overall F-scores. To make sense of overall F-scores, however, we need a function from the space of F-values to a one-dimensional scale. To denote the associated property associated with that scale—which we might call overall F-ness—I'll use the convention F°. Alongside that, it's useful to define a few other concepts:

the **F-scale**	a one-dimensional space of F-scores
an **F-score**	an aggregation of values along the dimensions of F, yielding an overall F-ness value[43]
F°	the property whose values are F-scores

The relationships between these concepts mirror the relationships between F-space, F-values, and F. These relationships are illustrated in the following diagram:

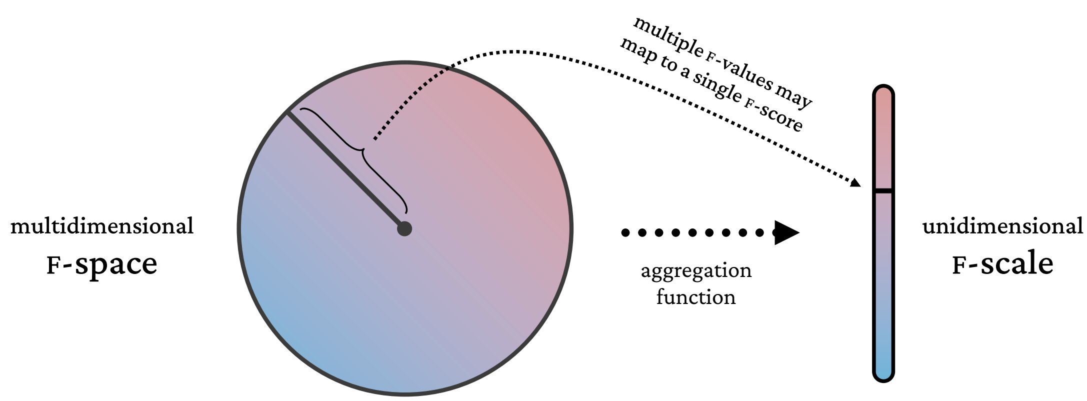

FIGURE 7: F VS. F°.

Now we are in position to identify a deeper relationship between dimensions and parameters. Consider a multidimensional property F whose values aggregate to scores along property F° (the F-scale). Let D be a dimension of F. What is the relationship that D stands in to F°?

[43] Notice that these definitions leave open whether overall F-scores must be interpreted as degrees of F. While many philosophers have assumed that aggregation over parameters of F requires taking overall F-scores to be degrees of F, it strikes me as unobvious whether we must interpret overall F-scores in that way. In my view, it's better to leave these definitions neutral on what kind of structure is built into the notion of an overall F-score.

The answer: D is a ***parameter*** of F°! To illustrate with another example, consider athleticism.[44] Suppose that speed, strength, and skill are each dimensions of athleticism, and that these aggregate to yield an overall athleticism score. Then there's athleticism°, the unidimensional property whose values are athleticism scores. And what relation do speed, strength, and skill stand in to athleticism°? The answer is they are each parameters of athleticism°.

Once we distinguish F from F°, we are in position to see that dimensions and parameters aren't merely distinct concepts. The concepts are related in systematic ways, and an elegant picture develops when we identify those relationships. To express this in more general terms, we can formulate some structural principles connecting dimensions parameters, multidimensional properties, and unidimensional scales. Let P_F be a complete collection of parameters of F, and let P_i be an individual parameter from that collection. Then:

Dimensions | Parameters

D is a dimension of F	↔	D is a parameter of the F-scale.
P_i is a parameter of F°	↔	F-space is a quotient of P_F-space.

Now we have the conceptual resources to fully diagnose the Puzzling Observation from earlier. The puzzle, as we saw, arose from a conflation between dimensions and parameters. But now we can better understand why that conflation is tempting. There's oftentimes a systematic ambiguity between talking about F vs. F°, or F-space vs. the F-scale, or the space of F-values vs. the space of F-scores. This makes invocations of dimensions especially precarious, since the dimensions of F are precisely the parameters of F°. Given this, it's easy to overlook the distinction between dimensions and parameters. When thinking about aggregation, dimensions *are* parameters. But the subtlety is that the relevant kinds are different: the dimensions of F are parameters of F°.

The distinctions I've developed can be used to clarify a number of other philosophical issues. As another example, consider **comparability**. Let's say that F is *comparable* just in case for any x and y with F-values, either $x >_F y$ or $x <_F y$ or $x =_F y$.[45] A number of philosophers have thought that there's a connection between multidimensionality and comparability. But what exactly is the connection? The ideas developed in this section enable us to think about this issue more sharply.

If F is multiparametric, then there aren't any obvious consequences for comparability. After all, the mere fact that there are multiple parameters determining F-values leaves open the structure of F-space: the space of F-values could be totally ordered, or partially

[44] See, for example, D'Ambrosio & Hedden [2024], who use athleticism as an example of a multidimensional property.

[45] See Dorr, Nebel, & Zuehl [2023].

ordered, or even unordered. While sometimes multiple parameters aggregate to a unidimensional ordered space, other times the target space of the parameters will have a different structure.

However, if F is multidimensional—and if F has multiple degreed dimensions[46]—then it's plausible that the space of F-values will be at best partially ordered. This is because the intrinsic structure of F-space itself won't itself determine how to trade off along the dimensions of F to yield overall F-scores. Instead, we need a function from F-space to an F-scale. If the F-scale is itself totally ordered, then there's a sense in which F will satisfy comparability (or, more accurately, F° will satisfy comparability).

Another application concerns the semantics of **multidimensional adjectives**...

These observations don't solve all the puzzles concerning aggregation, comparability, or multidimensional adjectives. But I hope it's clear that the distinctions developed in this section enable us to approach those puzzles in a more systematic way.

These remarks express some initial comparisons between dimensions and parameters. But there's much more to explore here. My central point is merely that these concepts are distinct. By disentangling them, we can think more clearly about the philosophical questions that have been classified under the label 'multidimensionality', and we can better understand how dimensions relate to other concepts such as comparability and aggregation.

Conclusion

I've explored a number of questions about the dimension relation. At the heart of this paper is the quotient analysis: for D to be a dimension of F is for the space of D-values to be a quotient of the space of F-values.

The basic idea of a dimension—a way of varying with respect to a kind—is intuitive. Yet the concept, in my view, turns out to express a deep mathematical relation: namely, that of a quotient. This may initially feel surprising, since the mathematics of quotients can become complex and intricate. But while the local details vary from case to case, the common pattern—an operation involving abstraction and structure-preservation—turns out to be one of the most universal mathematical concepts.

The ideas in this paper are merely the start of a larger line of inquiry. I've mostly set aside questions about dimensionality and dimensionalization, and I've only briefly touched on the relationship between dimensions and other structural concepts, such as parts and parameters. But I hope I've illustrated how the subject-matter of dimensions is rich, fertile, exciting, and promising, with seemingly endless new questions to explore.

[46] Multidimensionality, by itself, doesn't entail incomparability. In order for comparability to even make sense, the relevant property must be degreed, and not all multidimensional properties are degreed (consider color). Furthermore, even degreed multidimensional properties can be comparable (consider velocity). Therefore, generating incomparability requires at minimum multiple degreed dimensions.

APPENDIX: Quotients and Partitions

In the paper, I defined quotients via structure-preserving maps. But I noted that quotients can also be defined via structure-inheriting partitions. This APPENDIX explains the latter way of thinking about quotients.

A quotient—on this approach—starts with a partition. The partition, when thinking about the dimension relation, is over the space of F-values. But a quotient not only divides that space into equivalence classes, but also specifies how the partitioned space inherits structure from its source space.

For most kinds of mathematical structures, only certain kinds of partitions count as quotients. The question is whether the partition "respects the structure of the space." In such cases, the equivalence relation that generates the partition is called a **congruence**. When a partition is generated by a congruence, the partition inherits structure from the source space. But what counts as a congruence depends on the kind of structure under consideration.

It's beyond the scope of this paper to explain congruences across all mathematical domains. Just as with the notion of a structure-preserving map, the exact definition varies across mathematical domains.[47] Because of this—and for the reasons expressed in §4—I'll focus merely on explaining how structure-inheritance works for **relational spaces**, meaning sets of elements endowed with some relations over those elements.

For relational spaces, structure-inheritance is defined by the "universal lift:" a cell stands in a relation just in case each element in a cell stands in that relation. To denote the cell of the partition that contains value F_i, I'll use '$[F_i]$'. Then, for any relation R, we can define $R([F_i], ..., [F_j])$ as holding iff $R(F_i, ..., F_j)$ for every combination of values across $[F_i], ..., [F_j]$.[48] In other words, a tuple of cells stands in a relation just in case every corresponding tuple of their elements stands in that relation.

To apply this to an example, consider again hue and color. Suppose we partition color space by hue, so that $F_1 \sim F_2$ iff F_1 and F_2 determine the same hue value. The resulting partition abstracts away from all differences in saturation and brightness, leaving intact only differences in hue. Each equivalence class in the partition will preserve relations between hue values while forgetting relations corresponding to brightness and saturation values. More precisely, equivalence classes of the partition will stand in the angular distances that characterize the hue circle.

---

[47] Here are a few prominent examples. For topological spaces, the partitioned space is equipped with the quotient topology, the finest topology that permits a continuous mapping from F-space to $F/\sim$. For partially ordered sets, structure-inheritance is defined by the "existential lift," where if $F_1 \geq F_2$, then $[F_1] \geq [F_2]$. For metric spaces, structure-inheritance is defined by equipping the partitioned space with the least-distance metric, where the distance between two cells is defined as the infimum of the distances between their elements.

[48] More formally, $R([F_1], ...,[F_n]) \quad \leftrightarrow \quad \forall F_1 \in [F_i] ... \forall F_n \in [F_n] \;\; R(F_1, ..., F_n)$.

There's a deep connection between quotient maps, quotient spaces, equivalence relations, and structured partitions. In brief, D-space is a quotient space of F-space just in case there's some equivalence relation ~ such that the structured partition F/~ is **isomorphic** to D. In other words, a quotient space of F is structurally equivalent to a structure-inheriting partition of F. This idea can be expressed symbolically as follows (≅ means 'is isomorphic to'):[49]

$$\text{D is a \textbf{quotient} of F} \quad \leftrightarrow \quad \exists(q : \text{F} \longrightarrow \text{D}) \ : \ (\text{F}/\sim_q \ \cong \ \text{D}),$$
$$\textit{where } \text{F}_1 \sim_q \text{F}_2 \ \leftrightarrow \ q(\text{F}_1) = q(\text{F}_2).$$

This elegant connection between quotient spaces and structured partitions justifies the two ways of defining quotients. If we care only about differences up to isomorphism, then quotient spaces and partitioned spaces are indeed distinct ways of looking at the same mathematical object.

[49] This is a generalization of the First Isomorphism Theorem, a theorem in abstract algebra. See Burris & Sankappanavar [1981] on the theorem, and Awodey [2010] on the categorical generalization of quotients.

References

Awodey, Steve. *Category Theory*, 2nd ed., Oxford Logic Guides 52, Oxford University Press, 2010. ISBN 978-0-19-923718-0.

Burris, Stanley & Sankappanavar, H. P. *A Course in Universal Algebra*. Graduate Texts in Mathematics 78, Springer-Verlag, 1981.

Calosi, Claudio & Wilson, Jessica (2021). Quantum indeterminacy and the double-slit experiment. Philosophical Studies 178 (10):3291-3317.

D'Ambrosio, Justin & Hedden, Brian (2024). Multidimensional Adjectives. Australasian Journal of Philosophy 102 (2):253-277.

D'Ambrosio, Justin & Stoljar, Daniel (ms). The Indeterminacy of Consciousness.

Dorr, Cian, Nebel, Jacob M. & Zuehl, Jake (2023). The Case for Comparability. Noûs 57 (2):414-453.

Ferreiros, J.; Gray, J. J. (2006-04-27). The Architecture of Modern Mathematics: Essays in History and Philosophy. OUP Oxford. ISBN 978-0-19-151379-4.

Funkhouser, Eric (2006). The determinable-determinate relation. Noûs 40 (3):548–569.

Green, E. J. (2020). The Perception-Cognition Border: A Case for Architectural Division. Philosophical Review 129 (3):323-393.

Hedden, Brian & Muñoz, Daniel (2024). Dimensions of Value. Noûs 58 (2):291-305.

Jacobs, Caspar (forthcoming). In Defence of Dimensions. British Journal for the Philosophy of Science.

Jalloh, Mahmoud (2023). Dimensional Analysis: Essays on the Metaphysics and Epistemology of Quantities. Dissertation, University of Southern California

Jacobson, Hilla (forthcoming). On the Very Idea of Valenced Perception. Journal of Philosophy.

Kennedy, C., & McNally, L. (2005). Scale Structure, Degree Modification, and the Semantics of Gradable Predicates. Language, 81.

Kirby, Kyle (forthcoming). The Aesthetic Dimension of Value. Australasian Journal of Philosophy.

Kulvicki, John (2020). Modeling the Meanings of Pictures: Depiction and the Philosophy of Language. Oxford, GB: Oxford University Press.

Lee, Andrew Y. (2021). Modeling Mental Qualities. *The Philosophical Review* 130 (2):263-209.

Mac Lane, Saunders (1998). Categories for the Working Mathematician. Graduate Texts in Mathematics. Vol. 5 (2nd ed.). Springer-Verlag. ISBN 0-387-98403-8. Zbl 0906.18001

Maley, Corey (2026). Structural representation is analog representation. Philosophy and the Mind Sciences 7 (1).

Mason, Sam (2025). Dimensions of Emotional Fit. The Philosophical Quarterly 75 (1):125-146.

Melamedoff, Ariel (forthcoming). Neutral Monism and the Attributes of God in Spinoza's Metaphysics. Archiv für Geschichte der Philosophie.

Richeson, D. S. (2021, September 13). The journey to define dimension. Quanta Magazine. Retrieved from https://www.quantamagazine.org/a-mathematicians-guided-tour-through-high-dimensions-20210913/

Sagan, Hans (1994), Space-Filling Curves, Universitext, Springer-Verlag, doi:10.1007/978-1-4612-0871-6, ISBN 0-387-94265-3, MR 1299533.

Stevens, S. S. (June 7, 1946). "On the Theory of Scales of Measurement" (PDF). Science. 103 (2684): 677–680.

Strawson, Galen. *Stuff, Quality, Structure*.

Varzi, Achille, "Mereology", The Stanford Encyclopedia of Philosophy (Spring 2019 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/spr2019/entries/mereology/>.

White, P. Quinn (2025). Beautiful, troubling art: in defense of non-summative judgment. Philosophical Studies 182 (3):775-799.

Zell E, Lesick TL. Big five personality traits and performance: A quantitative synthesis of 50+ meta-analyses. J Pers. 2022 Aug;90(4):559-573. doi: 10.1111/jopy.12683. Epub 2021 Nov 1. PMID: 34687041.