The Harmony of Experience

Michael Pelczar

Abstract

The concept of phenomenal harmony or coherence—the idea that experiences can "fit together" in a more or less orderly or realistic way—plays a central role in classic idealist theories, as well as more recent neo-Kantian metaphysics of appearance and reality. Despite its importance to these philosophical systems, the concept of phenomenal harmony has never received an adequately detailed explanation. This paper attempts to repair that omission.

1 Introduction

Take all the conscious experiences you've ever had. Some of them you believe to have been veridical perceptions; those are the experiences that you reach for to justify your beliefs about the external world. Others you don't take to have been veridical perceptions, but dreams, optical illusions, hallucinations, etc.; you don't look to those experiences when you want to justify your beliefs about the external world.

The former experiences cohere with the totality of your experiences in a way that the latter don't. Other things being equal, the mental life of someone who never had any dreams or hallucinations, and never fell prey to any illusions, would exhibit more phenomenal coherence than the mental lives of people who do dream, hallucinate, etc. The former mental life would be more harmonious, phenomenologically, than the latter.

The idea of phenomenal coherence or harmony is central to the broadly idealist metaphysics that dominated the philosophical scene from the late seventeenth century through the first decades of the twentieth. For Leibniz, the physical world is just a certain harmony among different monads' perceptions. For Berkeley, a physical object is a combination of ideas that fits into the totality of ideas that God gives His creatures in a harmonious way. For Kant, the existence of a physical thing reduces to the fact that a complete realization of noumenal power (to give minds sensible intuitions) would include intuitions that harmonized with all other intuitions in certain ways. For Mill, for physical objects to exist is just for it to be the case that sensations tend to occur in such a way that some combinations of sensations relate harmoniously to the totality of all sensations.[1]

In the idealist view, a veridical experience (that reveals the world to be as it really is, in some respect) isn't one that the physical world causes in the right way, but one that relates to the totality of all experiences, or potential experiences, in the right way. Veridical experiences are those that relate to the experiences in a hypothetical realization of all potential for experience in the way that the experiences you're having now relate to the rest of the experiences you've had, rather than in the way that the experiences you have when dreaming or hallucinating relate to the rest of your experiences.

Phenomenal harmony has a similar role to play in neo-Kantian metaphysics, where questions about the nature of what causes our experiences take a back seat to questions about the structure and phenomenal character of the experiences they cause.[2]

Phenomenal harmony is a familiar feature of our experience, but that doesn't make it easy to define or analyze. In this respect, it's far from unique. Most of us know what musical harmony is well enough to tell when it's present or absent, but few of us can give a satisfactory definition or analysis of it. If I ask you what it is about the voices of a Bach fugue that make them harmonize with one another, you probably couldn't say. Still, the voices do harmonize, and you know that this is due to the musical qualities of the individual voices and how they relate to one another in the fugue.

Likewise, if I ask you what it is about each experience in an ordinary waking stream of consciousness that makes it cohere with the rest, you probably couldn't say. Still, the experiences do cohere, and you know that this is due to the phenomenal qualities of the individual experiences and how they relate to one another in the stream.[3]

But it's one thing to acknowledge that we have a concept of phenomenal harmony, and another to make phenomenal harmony the foundation of an experienced-based metaphysics of the physical world. Before we put that kind of weight on the concept, it seems prudent to reassure ourselves that our grasp of it doesn't depend on a prior grasp of physical concepts; otherwise, we can't be confident that a non-circular reduction of physical facts to facts about harmoniously-related experiences is possible. The best way to reassure ourselves on this score would be to give a clear analysis of phenomenal harmony in purely non-physical terms. Such is the purpose of this paper.

[1]For Leibniz, see (Leibniz, 1712/2007, 249, 257) and (Leibniz, 1714/1989, 220); for Berkeley, see (Berkeley, 1710/1901, 257-58) (also (Robinson, 1994, 226-31), (Foster, 2000, 250-55), and (Foster, 2008, 107-113)); for Kant, see (Kant, 1781/1998, 295-326); for Mill, see (Mill, 1865/1979, 225-264) (also (Ayer, 1946-1947)). The concept of phenomenal harmony also plays a salient role in sense-datum theories; see, e.g., (Price, 1932, 204-272) and (Nicod, 1924).

[2]See esp. (Chalmers, 2010).

[3]I take it that this is what G.E. Moore is getting at when he says that although he has conclusive evidence that he is awake and not dreaming, he can't tell us what all his evidence is; see (Moore, 1939/1993, 169).

In §2, I explain how to describe various aspects of conscious experience in geometric terms. This will give us the tools we need to carry out our analysis of phenomenal harmony.

In §3, I define the notion of a spatially complete set of experiences (roughly: a set of experiences that we can piece together to form a geometrically complete object).

In §4, I define the notion of a perspectivally complete set of experiences (roughly: a set of experiences that captures every geometric perspective on an object).

In §5, I analyze phenomenal harmony in terms of spatially and perspectivally complete sets of experiences.

In §6, I conclude with some suggestions about how we might use the concept of phenomenal harmony, as analyzed here, in the service of a plausible and scientifically-informed metaphysics.

2 Phenomenal geometry

Suppose a group of us are standing in a desert in broad daylight, facing each other in a circle about fifty feet across. With one exception, we all see—or have visual experiences that prompt us to claim to see—a large saguaro cactus in the middle of the circle. The exception is Oliver, who says he sees nothing but sand where we claim to see a cactus. Puzzled, he takes a step forward, and (as he would put it) a cactus appears. Stepping back again, the cactus seems to disappear. Stepping back once more, it seems to reappear.

Each of us takes a turn standing where Oliver began, and each of us has the same experience: while standing at that particular spot, and that spot only, we have no visual appearance of a cactus, but only an empty patch of sand. Stranger yet, when standing at that spot, one's experience doesn't even give an indirect hint of the presence of a cactus. From that spot, the sand at the center of the circle doesn't appear indented or disturbed; there's no shadow like that which a saguaro cactus casts on a sunny day; birds that appear from other spots to fly around the cactus appear from the "Oliver spot" to fly straight through the space that the cactus occupies (or seems to occupy from all other vantage points); etc. It's as though anyone who stands at the Oliver spot occupies a parallel universe, identical to our universe save for the absence of the cactus.

What are we to make of all this? Do we conclude that we are confronted with a very unusual cactus that has no physical presence relative to a certain observation point? Or do we conclude that there's something about the spot where Oliver stood that produces an illusion of an empty patch of sand for anyone who stands there?

I think we would definitely draw the latter conclusion. At least, we would conclude that *someone* was subject to illusion—if not those who stand where Oliver stood, then everyone besides those who stand there (maybe the Oliver spot is the only illusion-free zone in this scenario).

Why is this?

It's because we take it to be in the nature of physical objects that their existence is not tied to a single standpoint, or to any proper subset of the totality of all standpoints. Of course there are physical objects that we can *perceive* from some but not all standpoints (due to distance, or an obstruction, or some other factor). But the idea of a physical object that *exists* from some but not all standpoints makes no sense.

Oliver's experience (empty patch of sand, etc.) doesn't fit into the totality that includes all the experiences of the people in the cactus scenario. I suggest that this lack of fit comes down to the fact that the inclusion of Oliver's experience prevents us from construing the experiences in the totality as different perspectives on a unique geometric whole.

Conscious visual experiences may not occur in physical space, but even if they don't, they lend themselves to description in geometric terms. In this regard, they're like various other things, such as numbers, that we can describe in geometric terms without presupposing that they exist in physical space.

When you look at Fig. 1, there is a salient phenomenal respect in which your visual image of (A) duplicates your visual image of (B), but does not duplicate

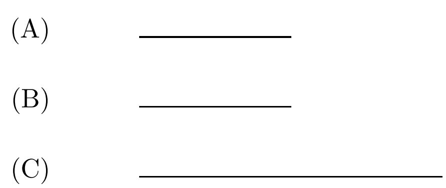

Figure 1: Phenomenal length

your visual image of (C); call it "phenomenal length." When you look at Fig. 2, your visual image of (D) resembles your visual image of (E) in a salient respect in which neither (D) nor (E) resembles (F); call it "phenomenal triangularity." When you look at Fig. 3, there's a salient phenomenal respect in which your visual image of (G) duplicates your visual image of (H), but not your visual image of (I); call it "phenomenal trihedronality"; when you look at Fig. 4, there's a salient phenomenal respect in which your visual image of (J) resembles your visual image of (K) but not your visual image of (L); call it "phenomenal cubicality."[4]

We can grasp phenomenal length, phenomenal triangularity, phenomenal trihedronality, phenomenal cubicality, and similar phenomenal qualities without thinking of the experiences that have them in physical terms. As we've just seen,

[4]Figs. 3 and 4 are a compromise. It would be better if I could present you with collections of three-dimensional objects like those that the images in the figures are meant to depict. Strictly speaking, I want to talk about the visual experiences you'd have if you were looking at such objects. Please work with me on this.

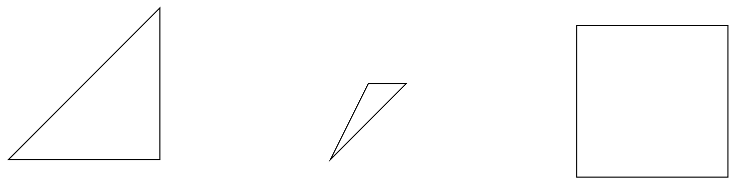

(D) (E) (F)

Figure 2: Phenomenal triangularity

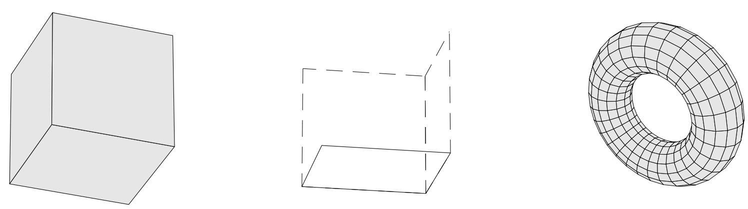

(G) (H) (I)

Figure 3: Phenomenal trihedronality

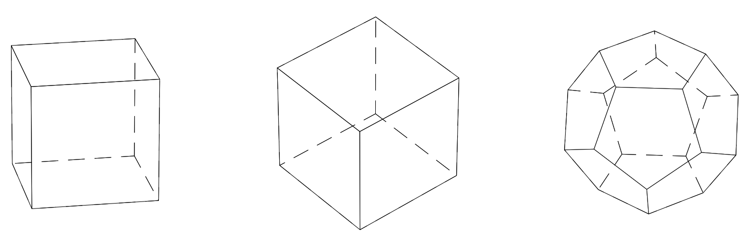

(J) (K) (L)

Figure 4: Phenomenal cubicality

we can grasp the qualities by simply reflecting on phenomenal similarities and differences found among our own experiences. There's no need think of the experiences as having physical properties, or as arising from physical events, or as representing physical objects. It's enough to notice that the experiences have certain phenomenal features.

The cases of phenomenal cubicality and trihedronality call for special comment. In my view, phenomenal cubicality, phenomenal trihedronality, and all other forms of phenomenal three-dimensionality are genuine phenomenal properties on a par with other phenomenal properties, like phenomenal squareness, phenomenal redness, and phenomenal loudness. In this, I agree with Price and Armstrong, and disagree with Berkeley, who basically denies that there is such a thing as phenomenal three-dimensionality, insisting that our visual experiences have only two-dimensional phenomenology, a third phenomenological dimension being, according to him, inferred from this two-dimensional phenomenology together with certain collateral non-visual phenomenology. It seems to me that Price and Armstrong are clearly right in this matter, and Berkeley clearly wrong, and I assume as much in what follows.[5]

Besides phenomenal shape and size, there are other phenomenal properties that lend themselves to description in geometric terms. For example, we may speak of phenomenal orientation, or phenomenal distance (or proximity).

One concept that will be especially useful in what follows is that of a *proximal phenomenal surface*.

Consider Fig. 5. Your conscious visual images of M and N are congruent with respect to total phenomenal surface. However, the total phenomenal surface of your visual image of O is not congruent with the total phenomenal surface of your image of M or N; the total phenomenal surface of the O-experience comprises six (rather than three) phenomenal squares.

You can see O either as a transparent cube that faces upward and to the left, or as a transparent cube that faces downward and to the right; this is the familiar Necker Cube gestalt-shift.

When you see O as an upward-facing cube, you see it as being such that the faces that meet at the point labeled 5 lie in front of the other three faces (so that if the faces meeting at vertex 5 were opaque, you couldn't see the other three faces). These three faces—the ones meeting at 5—correspond to the proximal phenomenal surface of the experience you have when you see O as an upward-facing cube.

When you see O as a downward-facing cube, you see it as being such that the faces that meet at the point labeled 3 lie in front of the other three faces (so that if the faces meeting at point 3 were opaque, you couldn't see the other three faces). These three faces—the ones meeting at 3—correspond to the proximal phenomenal surface of the experience you have when you see O as an downward-facing cube.

[5] See (Price, 1932, 244-45), (Armstrong, 1960, 65-73), and (Berkeley, 1709/1901, 127-28).

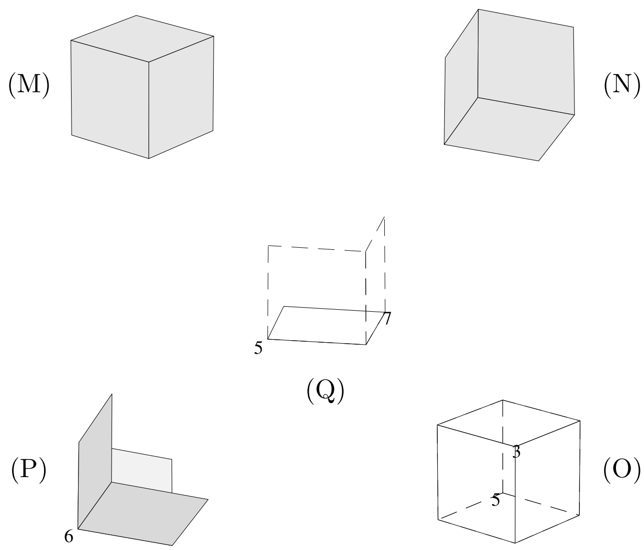

Fig. 5: Phenomenal surface

While the total phenomenal surface of the experience you have when you look at O differs from the total phenomenal surface of the experience you have when you look at M or N, the *proximal* phenomenal surface of the experience you have when you look at O (whichever experience it is, "upward" or "downward") is the same as the total phenomenal surface of the experience you have when you look at M or N. (For M and N, the total phenomenal surface just is the proximal phenomenal surface.) However, the proximal phenomenal surface of the experience you have when you look at P is not congruent with the total phenomenal surface of your M- or N-experience; the same goes for the experience you have when you look at Q (whether you see Q as facing down and to the left, or up and to the right).

The phenomenal geometry of a visual experience is typically more complicated that what you can describe just in terms of phenomenal lines, phenomenal cubes, or other phenomenal shapes that have official mathematical names. When I see a hedgerow, my experience has a nameless phenomenal geometry that I can best describe as "hedgerowish." It can also happen that one has an experience with a disconnected geometry. Such might be the experience of an astronaut who sees the Moon and Earth suspended in empty space. And of course phenomenal geometry isn't exclusive to visual experience: tactual experience and (to a lesser extent, perhaps) auditory experience can also have phenomenal-geometric properties. However, we'll concentrate on visual phenomenology here.

Instead of speaking directly of the phenomenal-geometric properties of our experiences, we can speak of such properties indirectly, using the following device.

For any given experience that has phenomenal-geometric properties, there is an abstract object whose geometry is congruent with the phenomenal geometry of that experience. For example, there's an abstract cube—a cube in an abstract Euclidean or quasi-Euclidean space—whose geometry is congruent with the phenomenal geometry of your K-experience.

Let's call an abstract object whose geometry is identical with the phenomenal geometry of a given experience ϵ an *abstract geometric projection of* ϵ*,* or "projection of ϵ," for short. So, instead of saying that your J-experience has the same phenomenal shape and size as your K-experience, we can say that your J- and K-experiences have congruent projections.[6]

I said that for any experience with phenomenal-geometric properties, there's an abstract object whose geometry is congruent with the phenomenal geometry of that experience. One might challenge this assertion, on the grounds that the phenomenal geometry of a visual experience is often, or even always, indeterminate in a way that the geometry of an abstract mathematical construct is not.

[6]In adopting this terminological device, we make no assumptions about the metaphysical status of abstract geometric objects. Maybe an abstract cube is a Platonic form. Maybe it's a mathematical function. Maybe it's a set of sequences of numbers. Maybe it's a surface of an abstract hypercube. Maybe it's an idea in the mind of God. For present purposes, it doesn't matter.

To assess this challenge properly, we need separate it from a different, and ultimately non-threatening, challenge to the determinacy of phenomenal geometry. When it comes to humans, at least, visual experience has only a finite phenomenal resolution. I can't tell just by looking whether the corner of a board is cut at ninety degrees, or at an angle that is greater or less than ninety degrees by a billionth of a degree. Nor can I tell just by looking whether two sticks are the same length, or differ in length by a billionth of an inch. However, none of this implies that there is anything indeterminate about the phenomenal geometry of my visual experience. It just implies that geometrically distinct physical objects sometimes give us geometrically identical visual experiences.

The real challenge to the claim that ordinary visual experiences have abstract projections comes from the fact that something can instantiate a phenomenal-geometric property ϕ without instantiating phenomenal counterparts of all the abstract geometric properties that must be instantiated in order for there to be an instantiation of an abstract counterpart of ϕ.

For example, when I look at a drawing of a hundred-sided polygon, my visual experience has the property of phenomenal many-sidedness. But there's no n such that my experience has the property of phenomenal n-sidedness. Yet, in order for an abstract object to have the geometric property of many-sidedness, there must be some n such that the object has the property of n-sidedness. Therefore (the argument goes) there's no abstract object whose geometry is congruent with the phenomenal geometry of the experience I have when I see a drawing of a hundred-sided polygon.

There are two ways to respond to this.

One way is to deny that there's no n such that my experience of the polygon has the property of phenomenal n-sidedness. True, there's no n such that I can tell by introspection whether my experience has the property of n-sidedness, but maybe that just shows that my experience has a phenomenal property that I'm unable to introspect.

If you're satisfied with this response, very well. Personally, I'm skeptical of the idea that we can have experiences with phenomenal qualities that are inaccessible to introspection (by us). An alternative response to the challenge at hand is to concede the point, and refine the notion of experiential projection accordingly. Instead of defining a projection of an experience ϵ as an abstract object whose geometry is identical with the phenomenal geometry of ϵ, we can define a projection of ϵ as the class of all abstract objects each of which has all of the determinate geometric features of ϵ.

So, for example, a projection of the experience I have when I look at a drawing of a hundred-sided polygon will be a set of abstract polygons. It won't be the set of *all* polygons. For example, it won't be a set that includes any triangles, since my experience has the property of phenomenal non-triangularity. But it will include 99-gons and 101-gons, since my experience doesn't have the property of non-99-gonity or non-101-gonity. (I say that this is true of my experience, not every possible experience. Maybe there are visual savants for whom there *is* such a thing as "looking like a hundred-sided figure," or "looking like a ninety-nine-sided figure"; maybe Stephen Wiltshire is such a person.)

At the end of the day, it doesn't really matter which definition of experiential projection we go with. Even though, or even if, the experience I have when I look at a hectogonal object has no property of phenomenal hundred-sidedness, I can count the sides of the object, and when I do, I have a stream of conscious experiences with determinate phenomenal geometries. Where necessary, we can use streams like that to construct hectogonal objects, rather than (humanly impossible) experiences of the sort that you'd need to be capable of having in order for there to be such a thing for you as "looking like a hectogon."

The same point applies to other phenomenal-geometric properties, like phenomenal length. Even if the experience I have when I look at a stick has no property of phenomenal 0.250000457-inches-longness, I could subject the stick to various scientific measurements, in the course of which I'd have a stream of conscious experiences with fully determinate phenomenal qualities. We can use such streams to construct objects with precise nano-scale lengths.

Just as we can speak of a projection of an experience into an abstract geometric space, we can speak of a projection of the proximal phenomenal surface of an experience into an abstract geometric space. Accordingly, let's distinguish between a *total projection* of an experience (what I have til now just been calling a projection) and a "surface projection" of an experience. A *surface projection* of an experience is an abstract geometric object that is congruent with the proximal phenomenal surface of that experience.

Surface projections are volumeless abstract geometric objects, geometric "films," so to speak. For example, a surface projection of the M, N, or O-experience is a geometric object that consists of three mutually perpendicular squares each of which shares two sides with the other squares.

A surface projection of the experience you have when you view P (with the most natural Gestalt) is a geometric object that consists of two mutually perpendicular squares that share a side, and a rectangle smaller than the squares that does not contact either square, but lies within a square region of space one corner of which is defined by the meeting of the squares at the point labeled 6.

On one Gestalt, a surface projection of the experience you have when you view Q is a geometric object that consists of two mutually perpendicular squares that share a side. On the other Gestalt, a surface projection of the Q-experience consists of three mutually perpendicular shapes: a square, a truncated square, and an irregular pentagon (the truncated square shares one vertex with the square, at point 7, and the irregular pentagon shares one vertex with the square, at point 5).

When you see an opaque object, a total projection of your experience is the same as a surface projection of the experience. As we've just seen, the same isn't true when it comes to transparent objects.

3 Spatial completeness

I now want to define a limited form of phenomenal coherence, in terms of experiential surface projections. This form of coherence is what I'll call "spatial completeness.” (Roughly but intuitively, a spatially complete set of experiences comprises experiences of the kind you’d expect a group of conscious observers to have, if they were surrounding some scene or object and viewing it from different angles.)

There are different senses in which a geometric object can be said to have sides. In one sense, a side of a geometric object is what we might call an “aspect” of the object. An *aspect* of an abstract geometric object O is the set S of all points on the surface of O that satisfies the following condition: there is a plane P that doesn’t contact O, such that for every point in S, the shortest line that connects that point to P does not pass through any point on the surface of O.

(To make this vivid, think of the illuminated surface of a sunlit object, and the rays of sunlight that connect points on that illuminated surface to a glass plate situated above the object and perpendicular to the Sun’s rays.)

Take a set E of experiences each of which has some phenomenal surface (not necessarily the same phenomenal surface for each experience). Now consider a set A that has as many members as E, and contains, for each experience in E, a surface projection of that experience.

There are infinitely many possible geometric combinations of the abstract objects in A. In some combinations, there is no contact among the objects; in some combinations, every object intersects or overlaps with some other object (or, as the case may be, with all the other objects); in some combinations, the objects are arranged roughly in a circle; etc.

Now ask whether, out of all these possible combinations, there is one combination—one abstract geometric object, G—such that (1) for every aspect of G, there’s an experience in E whose surface projection is congruent with that aspect of G, and, (2) for every experience in E, there’s an aspect of G that is congruent with a surface projection of that experience.

If the answer is “Yes,” then E is a spatially complete set of experiences.

Let me illustrate this with an analogy.

Suppose you have a small statue—say, a statue of a saguaro cactus. You also have a workshop set up to make thin plastic molds of objects such as statues. To make a mold, you start by setting the statue in front of a spray-nozzle. The sprayer coats the side of the statue facing it with a layer of resinous material. When the resin hardens, you pop it off of the statue; the result is a negative impression of one side of the statue. Now you put that pop-off in front of the nozzle, with the hollow side (the side that used to be in contact with the statue) facing the nozzle. The nozzle sprays that side of the pop-off with a rigid but extremely thin—in fact, infinitely thin—plastic coating. When the plastic hardens, you pop it off of the resinous impression. The result is an infinitely thin plastic object that perfectly matches the surface geometry of one side of the original statue.

Suppose you repeat this spray/pop/flip/spray/pop process many, many times. However, each time, you rotate the statue just barely (left, right, up, or down)—so little that for any given rotation, the statue appears to the naked eye not to have been rotated at all. You do this in such a way that by the time you’re done, the statue has been sprayed from every discernible angle, leaving you with a huge number of (infinitely thin) plastic casts of various sides of the statue.

This collection of casts has the following property: (1) each cast is congruent with some area of the surface of the statue, (2) each area of the surface of the statue (up to the limits of human visual discrimination) is congruent with some cast, and (3) there are just as many (discriminable, in most cases overlapping) sides of the statue as there are casts.

Suppose I walk into the warehouse where you store these casts, and find them stored in a big locker. I don't find any statues, and I don't have any prior knowledge of what the casts are, or where they came from. Having nothing better to do, I start playing around with them. I discover that because the casts are so thin and light, they can be made to stick to one another electrostatically just by putting them into contact. So I juxtapose them in various ways. I line them up in a row, to make a "scattered" object of non-contacting casts. I pile them in a heap to make a "gappy" object in which each cast contacts one or more other casts. Eventually, I figure out that I can make hollow cactus-statuettes out of the casts. I make an army of miniature cactuses and arrange them to look like a desert landscape (another scattered object).

In the course of this experimentation, I discover that the collection of casts has the following property: there is one, and only one, way of arranging them so that (1) the resulting object includes all the casts in the collection, and, (2) subjecting that object to the repeated spray/pop/flip/spray/pop process described earlier leaves you with a collection of casts indistinguishable from the collection with which you began (and out of which you assembled the object).

A spatially complete collection of experiences is analogous to a collection of casts that has the property just described.

These examples and analogies illustrate one way in which a set of experiences can be spatially complete, but I want to define spatial completeness somewhat more generally. The motivation for this is twofold.

First, it's possible for there to be a physical entity too large for any humanly conceivable experience to take in the entirety of each side of it. As it stands, our definition of spatial completeness doesn't allow there to be a spatially complete collection of experiences corresponding to a world that contains such an entity.

Second, it's not uncommon for a physical object to have a geometry such that not every part of the object's surface is included in an aspect of the entity, as we have defined "aspect." For example, as we've defined "aspect," various parts of the interior surface of an empty wine bottle don't belong to any aspect of the bottle (or of the bottle's counterpart in abstract geometric space), and no part of the interior surface of a tennis ball belongs to any aspect of the ball (or its abstract geometric counterpart). As it stands, our definition of spatial completeness doesn't clearly allow for there to be a spatially complete collection of experiences corresponding to a world that contains such objects.[7]

[7] I say "doesn't clearly allow," since we might try to deal with bottles and tennis balls by reducing them to collections of particles that have no hollows or involuted concavities.

Let's define a *partial aspect* of a geometric object O as the set S of all points on the surface of O that satisfies the following condition: there is a (possibly finite) plane-region P that doesn't contact O, such that for every point in S, the shortest line that connects that point to P does not pass through any point on the surface of O. (To make this vivid, think of the region of the surface of an object illuminated by a (possibly very small) spotlight. Now think of the rays of light that connect points on that illuminated region to a glass plate situated between the spotlight and the illuminated region, perpendicular to the spotlight's rays.)

Not every part of the surface of a tennis ball belongs to an aspect of the tennis ball, but every part of the surface of a tennis ball *does* belong to a *partial* aspect of the ball (or its abstract geometric counterpart). Likewise for the wine bottle.

Furthermore, even if there's an object too big for a human experience to take in any aspect of it, it's still possible to parse the object's surface into regions each of which is small enough for a human experience to take it in. Each such "small enough" region corresponds to a partial aspect of the object.

Here is the official definition of spatial completeness:

> A set of experiences E is *spatially complete* just in case there's an abstract geometric object G such that (1) for every partial aspect of G, γ, there's an experience in E whose surface projection is congruent with γ, and, (2) for every experience in E, ϵ, there's a partial aspect of G that is congruent with ϵ's surface projection.

4 Perspectival completeness

We are now almost in a position to give a geometric analysis of phenomenal harmony, but before we can do so, we need to introduce one last bit of technical apparatus.

Phenomenally distinct experiences can have the same phenomenal geometry. This can happen in various ways; for example, both a tactual and a visual experience might exhibit phenomenal flatness, and two visual experiences with different phenomenal colors might both exhibit phenomenal trihedronality.

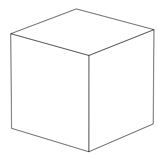

(R)

Fig. 6

Of particular interest to us is the possibility for experiences with the same phenomenal geometry to differ with respect to what we may call their "phenomenal perspective."

Your visual experiences of M and N in Fig. 5 illustrate this possibility. The images have the same phenomenal shape and size. Their surface projections are congruent. Yet, the experiences obviously differ from one another, in a phenomenal respect in which the M-experience doesn't differ from the experience you have when you look at the image (R) in Fig. 6.

Like phenomenal shape, size, and three-dimensionality, phenomenal perspective is a genuine phenomenal property on a par with other phenomenal properties, like phenomenal redness, loudness, and sweetness.[8]

Geometrically congruent experiences can also differ in terms of what we might call "phenomenal proximity." If you're looking at a barn from a distance of a hundred feet, then take ten steps back, your visual experience of the barn doesn't change in terms of its phenomenal geometry—it's a conscious appearance of a barn with the same size and shape—but it does change in some phenomenal respect, which is what I'm calling phenomenal proximity. For present purposes, it's convenient to classify phenomenal proximity as a kind of phenomenal perspective. (So, the two barn experiences just mentioned differ in phenomenal perspective, even if they are views of the barn from the same angle.)

Given any experience that exhibits phenomenal geometry, there is an *extendend perspective family* of experiences to which that experience belongs. An experience x belongs to the extended perspective family of an experience ϵ just in case (1) the surface projection of x is congruent with the surface projection of ϵ, and, (2) x and ϵ differ with respect to phenomenal perspective.

The extended perspective family of an experience includes many experiences that have the same phenomenal perspective. For example, the extended

[8]Here again I'm in agreement with Price: see (Price, 1932, 218).

perspective family of the N-experience includes both the M-experience and the R-experience. Define a *nuclear perspective family* of an experience ϵ as a subset of ϵ's extended perspective family that includes, for each member of ϵ's extended perspective family, one and only one experience with that member's phenomenal perspective. So, a nuclear perspective family—"nuclear family," for short—of the N-experience might include the M-experience or the R-experience, but not both.[9]

Let's say that a collection of experiences is *perspectivally complete* just in case it satisfies the following condition: for all x, if x belongs to the collection, then x has a nuclear perspective family all of whose members also belong to the collection.

And let's say that a collection of experiences is *worldlike* just in case it is both spatially and perspectivally complete.

Now back to Oliver and the cactus.

The collection of experiences that includes all the experiences described in the cactus case is not worldlike. For example, the collection doesn't include any visual phenomenology answering to the very top of the cactus (assuming that it's very tall), or to the subterranean parts of the cactus.

But suppose that we add more observers to the scenario (including, if necessary, avian and insect observers). If we add enough observers, and position them in the right way, the observers' experiences will constitute a spatially complete totality of experiences.

However, given the oddity of the scenario, the observers' experiences will not constitute a perspectivally complete totality.

We can imagine that Oliver stands shoulder-to-shoulder with two other observers, and that the visual experiences of those observers are geometrically congruent with the experiences that Oliver would be having if we were dealing with an ordinary cactus-viewing situation. So, the fact that no one who stands at the Oliver spot has a visual appearance of a cactus doesn't prevent the totality of experiences in this case from being spatially complete.

But it does prevent the totality from being perspectivally complete, since it means that the totality does not contain all the members of any nuclear family of the visual experiences of people standing in spots adjacent to the Oliver spot.

We can also imagine a situation in which there is no pathological Oliver-style experience, but several people are viewing a sphere suspended in empty space. (Suppose, for the sake of simplicity, that the people are invisible to one another.) Maybe there are six observers, looking at the sphere from above, below, the north, the east, the south, and the west, respectively.

The experiences of the people in this scenario also fail to constitute a worldlike collection, because their experiences don't constitute a spatially complete collection. (This is true, even if we make the collection perspectivally complete, by adding further observers at various distances from the sphere.) While it's true that one can assemble the surface-projections of their six experiences into

[9]Cf. Price's "gradual transition series" and "nuclear solids": (Price, 1932, 204-23).

a sphere—indeed, any two of their experiences is sufficient for that—one can't assemble the surface-projections into a sphere such that for every aspect of the sphere, there's a distinct surface-projection. In other words, there's no geometric object such that there is a one-to-one correlation between aspects of that object and surface-projections of experiences in the collection. Any sphere that you can assemble out of the six surface-projections will have many more than six discernible hemispheric surface regions.

One can assemble the surface-projections of the sphere-observers' experiences into a geometric object that has only six discernible hemispheric surface regions. The result is a hollow almost-hemisphere. But the aspects of this almost-hemisphere would include concave geometric forms, for which there are no congruent surface-projections among the surface-projections of the sphere-observers' experiences.

5 Phenomenal harmony

We can now give an analysis of phenomenal harmony. I'm going to do this by giving a preliminary analysis, which I'll then amend to deal with a problem related to phenomenal temporality.

Let's say that a set of experiences has *maximum phenomenal harmony* just in case it is worldlike (i.e., spatially and perspectivally complete).

Maximum phenomenal harmony is the highest grade of phenomenal harmony. A set of experiences has a lower grade, if the set isn't worldlike, but would become worldlike with the addition or subtraction of one or more experiences. The fewer total subtractions and additions required to achieve a worldlike set, the higher the original set's grade of phenomenal harmony.

The experiences in the cactus case constitute a set that has high, but not maximum, phenomenal harmony. The experiences that constitute a typical human mental life have a lower degree of harmony, but much higher than one would expect of an equinumerous set of randomly selected experiences.

Given a set of experiences, S, there is the set of all worldlike sets of experiences that differ minimally from S, i.e., that result from S by the smallest number of additions and subtractions. Call these the *minimal worldlike variants* of S. (If S is worldlike, S is its own minimal worldlike variant.)

An experience ϵ *phenomenally fits* into a set S of experiences to which ϵ belongs, just in case ϵ belongs to every minimal worldlike variant of S. (So, every member of a worldlike totality of experiences fits into that totality.)

Oliver's experience doesn't fit into the totality of experiences in the cactus scenario, since his experience doesn't belong to every (or any) mimimal worldlike variant of that totality. The fastest way to turn the totality of experiences in the cactus case into a worldlike totality is to replace the Oliver experience with a suitable visual appearance of a cactus (and add further experiences as indicated earlier: top of the cactus, cactus roots, etc).

Our preliminary analysis of phenomenal harmony is deficient in one important respect: it entails that *any* non-worldlike collection of experiences falls short of maximum phenomenal harmony. This is a problem, because, intuitively, we'd like to be able to count as maximally harmonious various collections of experiences that aren't worldlike. For example, it should be possible for the experiences of beings who take themselves (on the basis of their experience) to have experiences of a changing world over an extended period of time to qualify as a phenomenally harmonious totality of experiences. On the analysis proposed above, this isn't possible, since the totality isn't worldlike.

To remedy this defect, we can amend the proposed analysis as follows.

Take a given collection of experiences, C. Divide C into subgroups, each of which contains all (and only) the experiences in C that belong to the same stream of consciousness (so, one subgroup per complete stream of consciousness).

Now construct a table or matrix, such that each row of the matrix contains, for at least one subgroup, all the experiences in that subgroup, ordered contiguously in the row as they are ordered in the stream of consciousness in which they occur.[10]

There might be more than one matrix constructible from a given collection of experiences in compliance with these constructional rules. Also, it's consistent with the rules for a matrix constructed in compliance with them to contain rows with more than one stream of consciousness, or rows that contain blank cells (including possibly whole series of blank cells).

Suppose it's possible to put the experiences in a given set into a matrix of the sort just described, in such a way that the experiences in each column of the matrix constitute a worldlike totality of experiences. (We count an empty column as a degenerate case of a worldlike totality.) Call such a matrix of experiences a *worldlike matrix*.

A set of experiences has *maximum phenomenal harmony* just in case it's possible to arrange the experiences into a worldlike matrix. A collection of experiences has less-than-maximum phenomenal harmony according to how many additions or subtractions you'd have to make to the collection, in order to get a collection that has maximum harmony. The totality of actual human experience obviously falls short of maximum harmony. The most efficient way to convert the totality of actual human experiences into a worldlike totality involves throwing out dreams, hallucinations, etc. (and adding a host of possible experiences that nobody actually has). We can register this by saying that the experiences you have while dreaming or hallucinating don't fit into the totality of human experiences. (According to an idealist, a non-veridical experience just is one that doesn't fit into the totality of actual or potential experiences, in this sense of "fit.")

[10] I'm thinking of streams of consciousness along the lines developed in (Dainton, 2006), but other theories of the stream of consciousness are also consistent with the account of phenomenal harmony I'm developing here.

6 Conclusion

I've been focusing exclusively on visual phenomenology, but it would be easy to enrich our analysis of phenomenal harmony by bringing in other forms of experience that are also amenable to description in geometric terms, including tactile experience, and, to a more limited extent, auditory experience.

What about experiences or qualia that don't lend themselves to description in geometric terms? For example, what about olfactory experiences, or phenomenal color? According to the analysis given here, these aren't relevant to the harmoniousness or not of a collection of experiences.

This isn't a problem, given that the main purpose of analysing the concept of phenomenal harmony is to facilitate an experienced-based metaphysics of the physical world (whether Kantian, idealist, or phenomenalist). For that, it isn't necessary to define all forms of veridical experience in terms of phenomenal harmony; it's enough if we can define "veridical ϕ-type experience" in terms of phenomenal harmony, where ϕ-type experience is the only type we need in order to give an experience-based metaphysics of the physical world. So, as long as experiences as of shape, size, motion, and related geometric and chronometric phenomena are the only experiences we need to mention in order to carry out an experience-based metaphysics, a concept of phenomenal harmony is doing all the work we want it to do, provided that we can define the veridicality of such experiences in terms of it.

There is, of course, a difference between seeing a thing as having the color it actually has, and seeing it as having a different color (likewise for odor, flavor, temperature, etc). But, as the case of Frank Jackson's Mary illustrates, it seems possible for someone without color experience to learn everything about the physical nature of things, and the same presumably goes for someone without a sense of smell, taste, etc. Even though the usual way to learn about a thing's color is by means of phenomenally colorful experience, it's also possible to learn what color a thing has by having suitable achromatic experiences, such as achromatic experiences of the operations of a colorimeter or spectrophotometer. Likewise, although the usual way to learn about a thing's odor is by means of olfactory experience, it's also possible to learn what odor a thing has by having suitable non-olfactory experiences, such as visual experiences of an experiment using a gas chromatograph.[11]

None of this should be taken to imply that an idealist, Kantian, or phenomenalist must hold that the features that make a totality of experiences harmonious directly reflect the physical contents of reality.

In order for a physical world to have features corresponding to the phenomenal features that characterize a harmonious totality of human experiences, it would have to be more or less as classical, pre-twentieth (maybe pre-eighteenth) century conventional or scientific wisdom supposed. That's why the classic idealists, like Berkeley and Leibniz, thought that they could read-off

[11] For Jackson's case, see (Jackson, 1982).

the structure of the physical world from the structure of suitably well-behaved experience, with physical shape corresponding to phenomenal shape, physical duration corresponding to phenomenal duration, etc.

Today, we know that it's not that simple: even if what makes a totality of experiences harmonious is that the experiences in it have certain phenomenal-geometric properties, the physical facts that the totality entails need not (and usually do not) feature things with corresponding physical-geometric properties.[12]

The phenomenal geometry of our experience (epitomized by visual experience) is best described as that of a time and space in which objects have determinate shapes and sizes, processes have determinate durations, and there is such a thing as how things are arranged at any given moment. From a Newtonian standpoint, this also describes the basic geometry of the physical world.

The closest thing we find to this in the post-Newtonian world-view is the spacetime of general relativity. But in this spacetime, objects do not have determinate shapes or sizes, processes do not have determinate durations, and there is no such thing as how things are arranged at any given moment. The fundamental unit of distance in relativistic spacetime is the spacetime interval; in the spacetime of our experience, spatial and temporal distance are co-fundamental. In relativistic spacetime, every two points are connected by a path with spacetime length of zero; there is nothing analogous to this in the spacetime of our experience.

So, even though idealists, Kantians, and phenomenalists hold that the physical facts of our world supervene on the fact that the totality of all experiences (or potential experiences) includes veridical experiences as of objects with certain shapes and sizes, processes with certain durations, etc., they shouldn't hold that the physical features of our world have properties analogous to the phenomenal properties that make the subvenient totality of experiences harmonious.

The experimental physicist, like everyone else, has conscious experiences as of classical objects moving and changing in classical time and space. But considered as a whole, and in the limit of a full realization of phenomenological potential, these classical appearances may ground a physical reality that is distinctly non-classical. Our world contains long meetings around rectangular conference tables, but what it is for the meetings to be long and the tables rectangular isn't something we can read-off directly from the experiences we have of them, any more than we can read-off from our experiences the nature of color or odor.

[12] The need to distinguish the world's physical structure from the phenomenological structure on which it ostensibly supervenes was already known to Poincaré, Nicod, and Russell in the early 20th century; see (Poincaré, 1898), (Nicod, 1924, 187), (Russell, 1926), and Russell's preface to (Nicod, 1950).

References

Armstrong, David. 1960. *Berkeley's Theory of Vision*. Melbourne: Melbourne University Press.

Ayer, A.J. 1946-1947. Phenomenalism. *Proceedings of the Aristotelian Society*, **47**, 163–196.

Berkeley, George. 1709/1901. An Essay Towards a New Theory of Vision. *Pages 127–210 of:* Fraser, A.C. (ed), *The Works of George Berkeley*, vol. Vol. 1. Oxford: Clarendon Press.

Berkeley, George. 1710/1901. A Treatise Concerning the Principles of Human Knowledge. *Pages 233–347 of: The Works of George Berkeley, A.C. Fraser, ed.*, vol. Vol. 1. Oxford: Clarendon Press.

Chalmers, David. 2010. The Matrix as metaphysics. *Pages 455–494 of: The Character of Consciousness*. Oxford: Oxford University Press.

Dainton, Barry. 2006. *Stream of Consciousness: Unity and Continuity in Conscious Experience*. New York: Routledge.

Foster, John. 2000. *The Nature of Perception*. Oxford: Oxford University Press.

Foster, John. 2008. *A World for Us: The Case for Phenomenalistic Idealism*. Oxford: Oxford University Press.

Jackson, Frank. 1982. Epiphenomenal qualia. *The Philosophical Quarterly*, **32**(127), 23–36.

Kant, Immanuel. 1781/1998. *Critique of Pure Reason*. Cambridge: Cambridge University Press.

Leibniz, G.W. 1712/2007. Leibniz to Des Bosses, 16 June 1712. *Pages 255–257 of:* Look, Brandon C., & Rutherford, Donald (eds), *The Leibniz-Des Bosses Correspondence*. New Haven: Yale University Press.

Leibniz, G.W. 1714/1989. Monadology. *Pages 213–225 of:* Ariew, Roger, & Garber, Daniel (eds), *Philosophical Essays*. Indianapolis and Cambridge: Hackett Publishing Company.

Mill, John Stuart. 1865/1979. *An Examination of Sir William Hamilton's Philosophy, and of the Principal Philosophical Questions Discussed in his Writings*. Toronto: University of Toronto Press.

Moore, G.E. 1939/1993. Proof of an external world. *Pages 147–170 of:* Baldwin, T. (ed), *G.E. Moore: Selected Writings*. London: Routledge.

Nicod, Jean. 1924. *La Géometrie dans le Monde Sensible*. Paris: Librairie Felix Alcan.

Nicod, Jean. 1950. *Foundations of Geometry and Induction.* New York: The Humanities Press.

Poincaré, Henri. 1898. La mesure du temps. *Revue de Métaphysique et de Morale*, **6**, 1–13.

Price, H.H. 1932. *Perception.* London: Methuen & Co.

Robinson, Howard. 1994. *Perception.* London: Routledge.

Russell, Bertrand. 1926. The world of physics and the world of sense. *Pages 106–134 of: Our Knowledge of the External World as a Field for Scientific Method in Philosophy*, 2nd edn. London: George Allen & Unwin.