

Three Puzzles About Spatial Experience

David J. Chalmers

The three puzzles of the title concern three different aspects of spatial experience: orientation, size, and shape. For each aspect of experience, the puzzles ask whether a certain sort of lifelong perceptual illusion with respect to that aspect is possible.

Puzzle 1: Is it possible that everything is (and always has been) mirror-reversed: things that appear to be on your left are actually on your right, and vice versa?

Puzzle 2: Is it possible that everything is (and always has been) twice as big as it seems to be: things that appear to be one meter long are actually two meters long, and so on?

Puzzle 3: Is it possible that everything is (and always has been) stretched out in a certain direction: things that appear to be square are actually 2:1 rectangles, and so on?

I will argue for a (qualified) negative answer to all three questions. Illusions of all three sorts are impossible. This leads to conclusions about the content of spatial experience and spatial concepts: I will argue for a sort of spatial functionalism, on which space is picked out as whatever plays a certain functional role, over spatial primitivism, on which we have a more direct and primitive grasp of space. Finally, I suggest that this spatial functionalism leads to an anti-skeptical

⁰Forthcoming in A. Pautz and D. Stoljar (eds.) *Blockheads! Essays on Ned Block's Philosophy of Mind and Consciousness*, Oxford University Press. This paper combines a number of themes from Ned Block at the level of content and method. At the level of content: left-right reversal meets inverted earth. At the level of method: thought experiments about perceptual illusion and Twin Earth cases serve as a guide to the contents of perception and of thought. I am grateful to Ned for much enjoyable interaction over these and many other issues over the last two decades, and for his ongoing presence as a dear friend, colleague, and co-conspirator. I am also grateful to Brad Thompson for his obvious influence on this discussion, and to audiences at Arizona, Auckland, Bilkent, Brown, Birmingham, Bochum, Crete, Hertfordshire, Indiana, Oxford, Paris, Pittsburgh, Rice, Salvador, Singapore, Umea, and Western Australia. The discussion here is an elaboration of a discussion in Chapter 7 (section 5) of *Constructing the World*.

conclusion regarding many putative Cartesian skeptical scenarios, including brain-in-vat and Matrix scenarios: these scenarios do not involve systematic perceptual illusions.

1 First Puzzle: Left-Right Illusions

The first puzzle asks: is it possible that you have been under a lifelong left-right illusion: things that appear to be on your left are actually on your right, and vice versa? In such a mirror-reversal scenario, the house to your left will actually be to your right, cars that seem to drive on the left will actually drive on the right, and the hand that appears to be your left hand will actually be on your right side.¹

The sorts of possibility that are most relevant here is epistemic possibility. The question is roughly: can you be *certain* that things are not reversed in this way? Is it coherent to suppose that things are reversed in this way. Many people think that we cannot be certain that we are not brains in vats, and that it is coherent to suppose that we are in a Matrix scenario. It is roughly that standard that is at play here.

The sort of illusion that matters here is perceptual illusion: could things that perceptually appear to be on your left actually be on your right? Especially relevant is vision: what looks to be on your left is actually on your right. But the same issue arises for hearing, for touch, and for bodily awareness. Analogous issues also arise for belief and for language: could our beliefs about what is on our left be wrong in this way, or could our assertions be wrong? I will attend to some of those issues later, but for now the core issue concerns perception.

It is certainly possible that we could undergo *temporary* left-right illusions of this sort, at least in a single modality such as vision. There are left-right inverting lenses that mirror-reverse a retinal image so that things on one's left look to be on one's right. Of course the moment one acts or gets relevant information from the other senses, there will be conflicting information (one tries to move one's left arm and sees the arm on the right move), and perceptual adaptation of some sort will eventually ensue (after a few days, things on one's left will seem to be on one's left). But one will at least have visual left-right illusions at the start of this process.

One could imagine a more thoroughgoing temporary reversal. Perhaps one could combine the visual reversal with an auditory reversal (inverting headphones?) and a left-right flip of bodily

¹The discussion of the first puzzle, especially in the second half of this section, is considerably more intricate than the discussion of the second and third puzzles. Readers should feel free to skip to the second and third puzzles at any point and return to the first puzzle later.

inputs and of motor effectors. Then there would be no immediate sensory conflicts and the illusion would be harder to detect. In a familiar environment with left-right asymmetries, the illusion would be detected by violated expectations: one's house will be mirror-reversed, writing will go the wrong way, asymmetries in one's body will give things away. But if all this happened in an unfamiliar environment, and with a relatively symmetrical body, the illusion might in principle go undetected for a long time.

Could such an illusion last for a lifetime? One way to pose the question is as follows. Suppose we fitted a newborn baby with left-right inverting lenses, perhaps along with left-right inverters for other sensory inputs and for motor effectors. Would that person be subject to lifelong left-right illusions? The reader might pause at this point to contemplate a considered judgment about the case.

My own judgment is that this person would not be subject to lifelong left-right illusions. I will defer arguing for this conclusion. But to raise just one consideration in its favor: if we hold that a rewiring in sensory inputs like this can produce a lifelong perceptual illusion, we seem to open ourselves up to left-right skepticism, the view that we do not know what is on our left and what is on our right, and that we do not know which external direction is left and which is right. After all, it seems quite possible that evolution might have produced such a rewiring at some point (it is telling that the visual image on the retina is upside down). If it did, we would presumably suffer from left-right illusions. How do we know that there such not such a flip in our past, so that all of us are suffering lifelong left-right illusions? There seems no easy way to exclude this possibility. If we cannot exclude it, then left-right skepticism seems to follow.

Of course the argument from a yes answer to left-right skepticism is not entirely conclusive: various familiar anti-skeptical maneuvers (reliabilism, fallibilism, externalism about evidence, contextualism) might be used to resist it. Still, one might think that something has already gone wrong once this sort of skepticism has begun to threaten.

One can pose a number of closely related questions: Can there be inversions without illusions for left-right experience? Can there be Twin Earth cases for 'left' and 'right'? Are left and right directly presented in experiences thereof? To illustrate these questions and the options for answering them, it is useful to review some analogous questions in the more familiar domain of color and color experience.

Consider Ned Block's famous case of Inverted Earth. On Inverted Earth, everyone is fitted with color-inverting lenses, so that red light (that is, light in the wavelengths associated with redness) produces the retinal response that green light normally produces and vice versa. At the same time,

the environment has inverted colors on Inverted Earth: grass is (what we call) red for example, though Inverted Earthlings call it 'green'; blood is (what we call) green though Inverted Earthlings call it 'red'. As a result of these two inversions, grass on Inverted Earth produces exactly the same brain response as grass on Earth, and so on. In fact, we can suppose that the people on Inverted Earth have brains that are physically identical to their counterparts on Earth.

The question then arises: are the inhabitants of Inverted Earth undergoing color illusions? Are the inhabitants of Earth undergoing color illusions? Are both undergoing illusions? Or neither?

A number of views of this case are available. But one view seems most natural: Block argues for it in his original paper on Inverted Earth, and it seems to be by far the most common view in practice. This is the view that neither Earthlings nor Inverted Earthlings are suffering from an illusion. The word 'green' on Inverted Earth refers to (what we call) red, so when they say 'grass is green' they speak truly. When Inverted Earthlings see grass, they have experiences phenomenally identical to those we have when we see grass: we might say that these are greenish experiences. (I will adopt the convention throughout this paper of using 'X-ish' for the phenomenal property associated with experiences as of X in normal humans in the actual world. So 'greenish' is a predicate of experiences while 'green' is a predicate applied to external things.) But where greenish experiences in Earthlings represent greenness, greenish experiences in Inverted Earthlings represent redness. So when Inverted Earthlings see red grass and have a greenish experience in response, their experiences are veridical and not illusory.

This view is naturally combined with a view on which color experiences represent whatever color normally brings them, and on which colors are certain physical properties, such as surface reflectances. On Earth, reddish experiences are normally brought about by a certain reflectance r (the reflectance blood has on Earth), so these experiences represent r . On Inverted Earth, reddish experiences are normally brought about by a certain reflectance g (the reflectance grass has on Earth), so these experiences represent g .

We might say that a property is directly presented in an experience when necessarily, experiences of that phenomenal character represent that property. On the view just described, color experiences do not directly present color properties. Instead, the color property represented by an experience depends in part on non-experiential facts about the environment. In effect, our relation to color is somewhat less direct than it might have phenomenologically seemed.

The same goes at the level of language. On this view, Earthlings use the word 'red' to refer to r , whereas physically identical Inverted Earthlings use their word 'red' to refer to g . This behavior is structurally analogous to that of 'water' in Putnam's (1975) thought experiment involving Twin

Earth, a planet just like Earth except that H₂O is replaced by the superficially identical XYZ. Earthlings use ‘water’ to refer to H₂O, while Twin Earthlings use ‘water’ to refer to XYZ. Let us say that a word is *Twin-Earthable* if there is a Putnam-style Twin Earth case involving that expression: that is, if there is a possible speaker using the term (nondeferentially) with one referent and there is a possible duplicate speaker using a corresponding term with a different referent.² Then on these views, ‘water’ and ‘red’ are Twin-Earthable, in contrast to terms such as ‘zero’ and ‘person’ that are plausibly not Twin-Earthable.

We might call this view *color functionalism*, in that colors such as redness are picked out for us in virtue of their causal or functional role: in particular, their role in bringing about certain sorts of color experiences. This is not a view about the metaphysics of color: it is consistent with the view that colors are functional or dispositional properties (e.g., redness is the disposition to cause reddish experiences), but it is also consistent with the view that colors are physical properties. Rather, it is a view about how reference to colors is fixed. It can also be construed as a view about color concepts, on which these are functional concepts: our concept of redness is the concept of whatever property plays a certain role in causing color experiences.

Whether or not this is the correct view of color experience, I think that something like it is very plausibly the correct view of left-right experience.

The nearest analog of Block’s Inverted Earth is what we might call ‘Mirror Earth’. On Mirror Earth, everyone is fitted from birth with left-right-reversing devices for sensory inputs and also for motor effectors. At the same time, the environment is left-right inverted with respect to Earth. As a result, the brains of people on Mirror Earth may be physically identical or at least similar to those on Earth. There may be various differences: brain surgery (and other interactions between brain and environment not mediated by sensors and effectors) may lead to divergences from Earth, and asymmetries in biology and in physics may lead to other divergences. But the case is at least enough like Inverted Earth that we can raise similar questions.

(An easier-to-grasp variant on Mirror Earth is a world where brains are left-right reversed, and bodies and environment are normal, with sensorimotor connections rewired to compensate. Mirror

²The qualification ‘nondeferentially’ is used to set these Twin Earth cases aside from Burge-style Twin Earth cases that turn on semantic deference to a linguistic community. One can arguably construct a Burge-style twin case for any term whatsoever, even terms such as ‘zero’ for which Putnam-style cases seem impossible. As defined here, such terms are not Twin-Earthable. Duplicate speakers are best construed as functional and phenomenal duplicates rather than physical duplicates. For much more on Twin-Earthability, see excursus 18 in the extended version of *Constructing the World*.

Earth is in effect a left/right reversal of this variant.)

Are inhabitants of Mirror Earth undergoing left-right illusions? Are inhabitants of Earth? Are both? Or neither? I think there is a strong intuition (perhaps even stronger here than in the color case) that neither Earthlings nor Mirror Earthlings are undergoing an illusion. When a tree is on the right of a Mirror Earthling, they have the same sort of experience an Earthling has when a tree is on their left. For an Earthling, this experience represents the tree as being on the left. For a Mirror Earthling, this experience represents the tree as being on the right. So both experiences are veridical.

On this view, ‘left’ and ‘right’ are Twin-Earthable. ‘Left’ refers to left for an Earthling, but to right for a Mirror Earthling. Likewise, the distinctive sort of experience that Earthlings have when things are on their left—call these leftish experiences—represent things as being on the left for Earthlings, but represent things as being on their right for Mirror Earthlings. So the relation of being on the left is not directly presented in this sort of experience. Rather, it depends both on the experience and on the environmental circumstances.

We might call the resulting view *left-right functionalism*. On this view, left and right are picked out in virtue of the causal roles in causing our left-right experiences. At the level of language: the external relation we call ‘left’ is picked out as the relation l such that l normally causes leftish experiences (that is, such that things that stand in l to us normally cause leftish experiences). At the level of concepts: our concept of *left* is a concept of what normally causes leftish experiences in us. At the level of perception: leftish and rightish experiences represent things as standing in l and r to us respectively, where these are the respective normal causes of leftish and rightish experiences in us.

Left-right functionalism entails that there are no permanent left-right illusions. According to left-right functionalism, leftish experiences represent the property that normally causes them. It follows that leftish experiences cannot represent one property and be normally caused by another. But permanent left-right illusions requires leftish experience to represent one property (left) while being normally caused by another (right). So left-right functionalism is incompatible with permanent left-right illusions.

Of course left-right functionalism is not the only possible view of the case. The main alternatives are varieties of *left-right presentationalism*, on which leftish experiences directly present a certain relation l : all leftish experiences represent l , and in effect we directly grasp the relation of being on the left. These views are analogous to color presentationalism, on which reddish experiences directly present a certain color r : in effect, we directly grasp the property of being red in

experience.

We can spell out various versions of left-right presentationalism in part by analogy with familiar views in the color case. We can distinguish the views according to whether they are *realist* or *eliminativist*: is the property/relation represented by these experiences instantiated in the actual world? We can also distinguish them according to whether they are *internalist* or *externalist*: is the phenomenal character of the relevant experiences (and therefore the nature of the represented property) determined wholly by the intrinsic properties of the subject, or partly by the environment?

Internalist realism: On this view, the phenomenal character of the relevant experiences depends on the intrinsic properties of the subject, and all experiences with this phenomenal character represent a certain property, one that is actually instantiated in our environment. In the case of color, the most plausible version of this view is a form of color primitivism on which red is a primitive property instantiated in our world, whose nature is grasped directly in (internally grounded) reddish experiences.³ Likewise, in the case of left and right, the most plausible version of this view is a form of left-right primitivism on which left is a primitive relation instantiated in our world, whose nature is grasped directly in (internally grounded) leftish experiences.

Internalist realist color presentationalism leads naturally to a view on which reddish experiences in Earthlings are veridical (red apples are correctly represented as red), while reddish experiences in Inverted Earthlings are illusory (red apples are incorrectly represented as green). Other versions of the view are possible, including one on which only Inverted Earthlings get it right.

Likewise, internalist realist left-right presentationalism leads naturally to a view on which leftish experiences in Earthlings are veridical while leftish experiences in Mirror Earthlings are nonveridical. Other versions of the view are possible, including one on which only Mirror Earthlings get it right.

One can argue against this form of color presentationalism on the grounds that it leads to color skepticism: it would seem something of a happy accident if we get things right and Inverted Earthlings get things wrong, and there seems little reason to think that we are in the good case. One can also argue on the basis of science that it is implausible that these primitive color properties

³What about internalist realist forms of color physicalism and color dispositionalism? The former is made implausible as there seems to be nothing intrinsic to the subjects that indicates that reddish experiences should represent one reflectance rather than another. The most plausible form of the latter is one on which reddish experiences represent the disposition to cause reddish experiences in the subject of the experience; but now this will be different dispositions in different subjects, so it is not a form of color presentationalism and is best regarded as a form of color functionalism.

are instantiated in our world. The same arguments are available in the case of left and right. As we saw earlier, this form of left-right presentationalism tends to lead to left-right skepticism: it would seem something of a happy accident if we get things right, and there seems little reason to think that we are in the good case. Furthermore, one can argue on the basis of science that it is implausible that primitive left-right properties are instantiated in our world. I will develop this last point shortly.

Externalist realism: On this view, the phenomenal character of relevant experiences depends on the environment, and experiences with that character represent a property instantiated in the external environment. In the case of color, the most plausible version of this view is one on which both the phenomenal character and the representational content of color experience depend on the external properties that normally cause those experiences: reddish experiences are normally caused by redness (which could be a primitive property or a physical property), and thereby represent redness and have the phenomenal character of reddishness. The same goes for the corresponding view of left and right: leftish experiences are normally caused by the left relation, and thereby represent left and have the phenomenal character of leftishness.

Externalist realist color presentationalism leads naturally to a view on which reddish experiences in Earthlings are veridical (red apples are correctly represented as red), while reddish experiences in Inverted Earthlings are also veridical (red grass is correctly represented as red). On this view, when an Earthling seeing an apple has a reddish experience, a duplicate Inverted Earthling seeing an apple will have a greenish experience, and both experiences will be veridical. Likewise, externalist realist left-right presentationalism leads naturally to a view on which leftish experiences in both Earthlings and Mirror Earthlings are veridical. When an Earthling seeing an object to the left has a leftish experience, a duplicate Mirror Earthling seeing an object to the right will have a rightish experience, and both experiences will be veridical.

One can argue against this form of color presentationalism on various grounds. One ground comes from cases of indeterminacy: if someone from Earth spends enough time on Inverted Earth, the content of their experience will gradually become indeterminate or divided between red and green, but it is hard to make sense of an experience whose phenomenal character is indeterminate or divided between reddish and greenish. One can also argue on the grounds of science that primitive color properties are not instantiated, and representing physical properties cannot yield a substantial phenomenal difference between reddishness and greenishness. Related arguments are also available in the case of left and right. It is hard to make sense of an experience whose phenomenal character is indeterminate or divided between leftish and rightish. Science suggests

that primitive left-right relations are not instantiated, and the relations that are instantiated could not ground a substantial phenomenal difference between leftishness and rightishness. Again, I will develop this last point further shortly.

Eliminativism: On this view, experiences of the relevant sort represent a property that is not instantiated in the external environment. In the case of color, the most plausible version is a sort of primitivism on which reddish experiences represent a primitive color property that is not instantiated in our world: one can think of this as Edenic redness, a pure qualitative property that might have been instantiated in the Garden of Eden but is not instantiated in the world revealed by science. In the case of orientation, the corresponding view is one on a view on which leftish experiences represent Edenic leftness, a primitive relation that might have been instantiated in the Garden of Eden but is not instantiated in the world revealed by science.

On this view, color experiences on Earth and Inverted Earth will all be illusory, as will left-right experiences on Earth and Mirror Earth. This seems sufficient reason for many to reject the view in the case of color: it seems a great cost to many to allow that so many experiences are illusory and that apples are not really red. The same applies even more strongly in the case of orientation: it is even harder to allow that nothing is really on our left. Furthermore, while there may be some phenomenological motivation for the view in the case of color, this phenomenological motivation is weaker in the case of orientation, for reasons I discuss shortly.

Once we have ruled out these forms of color and orientation presentationalism, color and orientation functionalism are left on the table as the most plausible views.

Categoricalism and relationism about orientation

To fully rule out orientation presentationalist views, there are two points from the discussion of presentationalism that we need to attend to further. First, whether there are primitive relations of left and right in the physical world. Second, whether there is a deep categorical difference between the experiences of left and right.

On the first issue: The primitivist about left and right holds that left and right are distinct “absolute” primitive orientations that we grasp in experience. The idea here is that they orientations are not barely distinct from one another. Rather, they are categorically distinct: each of them has its own categorical nature in virtue of which they are distinct. Insofar as the primitivist is a realist, she will hold that left and right are instantiated in the physical world. So the physical world will involve a categorical difference between left and right orientations.

The idea of a categorical difference between orientations may seem elusive, but we can pin it down in a familiar way. We can ask: is there a possible world, distinct from this one, that is

a physical duplicate of this world except that it is left/right reversed? If there is a categorical difference between left and right in the physical world, we would expect this to be possible. If there is no such difference, we would expect this to be impossible: a putative left/right reversed version of this world would be physically identical to our world.

Henceforth I will call *physical orientation categoricism* (or just *p-categoricism*) the thesis that a world might differ physically from ours merely in being mirror-reversed, and *physical orientation relationism* (or just *p-relationism*) the thesis that a world could not differ physically from ours in this way. As is often the case with modal definitions (e.g. definitions of physicalism in terms of supervenience), this definition may not precisely characterize the most fundamental underlying issue, but it gets at a relatively straightforward thesis in the near vicinity.

The most famous p-categoricist was Immanuel Kant. In “On the first ground of the distinction of regions in space” (1768), Kant considers incongruent physical counterparts, such as left gloves and right gloves, and argues that there is a difference between them that goes deeper than the mere relational fact of their incongruence (the fact, roughly, that one cannot be moved onto the other). He suggests that even in a world that contains only a single glove, there would be a fact about whether it is a left glove or a right glove. In effect, he suggests that there are two distinct mirror-reversed worlds, one with a left glove and one with a right glove. So Kant advocates p-categoricism at least about those worlds. What goes for those world presumably also goes for our own more complex world, in which case his view is also p-categoricist about our world.⁴

Despite Kant’s argument, the consensus among contemporary philosophers of physics (Hofer 200, Pooley 2003, Baker 2011) is that p-categoricism is false. Left gloves and right gloves are distinct merely in virtue of being incongruent, and not in virtue of any underlying absolute property. There are not distinct worlds with a left glove and a right glove: there is just one world described two different ways. Likewise, a putative mirror-reversed version of our world would just be our world, described differently. It is not that p-categoricism is held to be incoherent or incompatible with the evidence. Rather, it is held that all the data are explained by p-relationism, so that further categorical differences explain nothing.

One might wonder here about parity violations in physics, where it appears that certain fundamental laws of nature involve a left/right asymmetry: roughly, certain charged particles decay

⁴The distinction between p-categoricism and p-relationism is independent of the more familiar distinction between relationism and substantivalism about space. A relationist about space could invoke absolute or non-absolute orientation relations; and a substantivalist about space could accept or deny that a world where both space and its contents are flipped is distinct from our own. Thanks to Geoff Lee for discussion here.

to the left rather than to the right. But even here, there is no categorical difference. One can still consistently hold that a mirror-reversed world would be the same world. As Pooley (2003) and Baker (2011) note, physics seems to recognize only one possible world in which particles decay in one direction and not the other, rather than two distinct worlds where particles decay to the left and to the right.

So there is good reason to think that p-categoricalism is false. There are two orientations in nature, left and right, but these are simply distinct orientations, not in virtue of any underlying categorical difference between the two. This view problematizes any realist view on which we grasp the distinct natures of physical left and physical right in left/right experience: it appears that there are no distinct natures to be grasped here. In particular, it problematizes realist left/right primitivism. If that view were right, mirror-reversal would certainly be possible, so its impossibility rules out this form of primitivism.

On the second issue: we can ask an analogous question about the *experience* of left and right. Is there an absolute categorical difference between leftish and rightish experiences? The experiential left/right categoricalist, or *e-categoricalist*, holds that there is such a difference: roughly, leftish experiences are intrinsically “leftish”, and rightish experiences are intrinsically “rightish”. The experiential orientation relationist, or *e-relationist*, holds that there is no such categorical difference. Leftish and rightish experiences are simply different experiences, with any apparent “qualitative” differences between them arising from differences in associations, memories, dispositions, and the like.

As before the question is somewhat obscure, but as before we can clarify it by asking: for a given subject with a total experience that is left/right asymmetrical, could there be a different subject with a mirror-reversed total experience that is phenomenally distinct from that of the first subject? For example, could there be a subject who has mirror-reversed visual fields and other sensory fields with respect to me, and in which all associations are reversed (he reads from right to left, drives on the right, and is left-handed), and whose total experience is phenomenally different from mine? The e-categoricalist says yes, while the e-relationist says no. As in the case of p-categoricalism and p-relationism, we can take this as our working definition of e-categoricalism and e-relationism.

An analog of e-categoricalism is extremely plausible in the case of color. There seems a categorical phenomenal difference between reddishness and greenishness, not a bare relational difference. This categorical difference is brought out by the apparent conceivability and possibility of inverted subjects in which reddishness and greenishness are systematically reversed. But the

corresponding categorical difference is at least much more elusive in the case of leftishness and rightishness.

My own tentative view is that e-relationism is true: there is no categorical difference between leftishness and rightishness. The issues here are intricate and their discussion is long, out of proportion to the small role that e-relationism plays in the rest of this paper, so I have left the discussion for an appendix at the end of this article.

How does the status of p-categoricism and e-categoricism affect the various views of left/right experience? I think as follows. The most natural forms of presentationalism seems to require e-categoricism, in order for us to have an experiential grip on left and right. It follows that if e-relationism is true, we should reject these versions of presentationalism. In addition, these forms of presentationalism seems to require p-categoricism in addition, in order that there to be properties for us to grip. So if we accept p-relationism, we should reject these realist forms of presentationalism.

There is perhaps room for version of presentationalist realism that is both e-relationist and p-relationist. For example, if there are nonabsolute leftness and rightness, there could also be corresponding nonabsolute leftishness and rightishness, consisting in awareness of these. On an externalist version, instances of leftishness might be grounded in relations to external instances of leftness, while on an internalist version they might be grounded in underlying internal instances of leftness. Views of this sort would not be subject to the objections from the falsity of categoricism, although they would still be subject to the other objections to internalist and externalist realism discussed above.

If e-categoricism and p-relationism are true, then one could in principle accept an eliminativist form of presentationalism: leftish experiences represent a primitive relation (Edenic leftness?) that nothing actually instantiates. But such a view is unattractive in its implication that orientation experience is always nonveridical and that claims such as 'X is on the left' are always untrue. A better reaction would be along the lines of the two-stage Edenic view of color experience in Chalmers (2006), which reacts to a analogous combination of views by saying that perfect veridicality of reddish experiences requires Edenic redness, a property not instantiated in our world, but imperfect veridicality requires only the physical properties that normally cause reddish experiences. Color terms such as 'red' refer to the latter properties, picked out in virtue of their causal role. In the case of orientation, we could say that perfect veridicality of leftish experience would require Edenic leftness, a relation not instantiated in our world, but imperfect veridicality requires only those physical relations that normally cause leftish experiences. Orientation terms

such as 'left' refer to the latter relations, picked out in virtue of their causal role. In this way, we will still have a sort of left/right functionalism.

If all this is right, then p-relationism, an orthodox view in physics, leads naturally to left/right functionalism. E-relationism also leads there. I have argued for both views here. So I think there is a good case for left/right functionalism.

We have seen that left/right functionalism entails that there are no permanent left/right illusions are metaphysically impossible. I have argued that left/right functionalism is correct, and I conclude that there are no permanent left/right illusions.

Admittedly, the arguments do not allow me to be certain that left/right functionalism is correct: for example, the argument for p-relationism is based on empirical and abductive considerations that fall short of certainty. Consequently, the arguments do not establish that left/right illusions are epistemically impossible, in the sense that I can rule out such illusions with certainty. For example, I cannot conclusively exclude the possibility that p-categoricism and e-categoricism are both true and that my experiences are systematically mistaken. I suspect that more ideal phenomenological reflection might conclusively reveal the truth or falsity of e-categoricism, which might allow a stronger conclusion. For now, I content myself with the observation that there is a good reason to believe left/right functionalism, and that I can know conclusively that *if* left/right functionalism is true, then permanent left/right illusions are ruled out.

Second Puzzle: Size Illusions?

The second puzzle asks: is it possible that you have been under a lifelong size illusion where everything is twice as big as it seems to be: things that appear to be one meter long are actually two meters long, and so on? If I am under such an illusion, then I am actually twelve feet tall rather than six feet tall, Olympics swimming pools are actually be 100 meter long rather than 50, cricket pitches will be 44 yards long rather than 22, and so on.

As in the first puzzle, the sort of possibility that is most relevant is epistemic possibility, and the sort of illusion that is most relevant is perceptual illusion. It is worth noting that one must be mistaken about everything in one's environment in this way: one cannot cheat by holding the size of one's body fixed, for example. More generally, there can be no trading on errors about relative size (the size of objects relative to one's body, for example). That is the domain of the third puzzle. For now, we are concerned with illusions about absolute size.

As before, it is certainly possible that I am under a temporary size illusion of this sort. It

is not so easy to devise size-doubling contact lenses, so instead we can appeal to body-doubling scenarios in the spirit of Alice in Wonderland. For now, suspend any doubts about the physical possibility of these scenarios; I will address that issue at the end of this section.

Suppose that overnight, while I am sleeping, my body doubles in size. In order that I do not notice immediately, we can suppose that my bed and my whole house double in size too. I will wake up and notice no difference. My hands seem normal size, as do my body and my bed. Looking in the mirror, my body seems to be about six feet tall, but this is an illusion: in fact it is twelve feet tall. I walk to the front door of my house and look outside. I exclaim with surprise: my car seems to be half the size that it used to be, the trees have shrunk, and so on. Again, this is an illusion: the car and the trees are exactly the same size they used to be.

As in the case of orientation, I will plausibly adapt to this situation before long and perhaps come to avoid size illusions. In fact there are versions of the scenario that may avoid illusions altogether. For example, perhaps I take a pill that I am told is a body-doubling pill, and I feel my body creak and stretch while I look out at a constant scene with trees and my car. In this case, it is not implausible that at the end of the process, the trees and my car will still look normal size to me, while my body will seem twice as large. (Perhaps one could test a version of this hypothesis using virtual reality devices.) All this brings out that size perception and size illusions are very much affected by background cues and background knowledge.

Still, in the first body-doubling case as described, I will at least have an initial stage of size illusion where things are twice as big as they seem. One could even imagine that this stage persists for a long time, if my whole environment over that period doubles in size. And it seems that I cannot exclude with certainty the possibility that I am in such a situation now. So temporary doubling illusions certainly are possible.

Could such an illusion last for a lifetime? One could imagine doubling the size of a newborn baby, for example, along with its whole environment. Would this lead to ongoing size illusions?

Perhaps the cleanest version of the case is the scenario that Brad Thompson calls Doubled Earth. Doubled Earth is just like Earth, except that everything is twice as big. (As before, doubts about physical possibility are deferred until later.) We can think of it as being in a distant part of our own universe. If we went there, we would find counterparts of us who are twice our size, engaged in physical activities that are otherwise very much like the activities we engage in on Earth.

Now consider Oscar on Earth, looking at a one-meter long ruler and having a normal visual experience of it, and his counterpart Doubled Oscar on Doubled Earth. Is Doubled Oscar undergoing

a size illusion? Is Oscar? Are both? Or neither?

In this case there is an overwhelmingly plausible verdict: neither of them is undergoing a size illusion (or at least, no size illusion deriving from the difference in size). Oscar's experience represents his stick as being one meter long, while Doubled Oscar's experience, this experience represents his stick as being two meters long. It is plausible that these two experiences are phenomenally identical, on which case we can conclude that absolute size is not directly presented in experience.

Something similar goes at the level of language. Oscar's expression 'one meter' picks out one meter, while Doubled Oscar's expression 'one meter' picks out two meters. It follows that 'one meter' is Twin-Earthable, at least if we count Doubled Oscar as a (functional) duplicate of Oscar. Something similar applies to the concepts that Oscar and Doubled Oscar express with these terms.

This verdict leads naturally to *size functionalism*. Linguistically: the size we call 'one meter' is picked out in virtue of the causal role it plays in causing size experiences. Conceptually: our concept of *one meter* is roughly a concept of what normally causes one-meter-ish experiences in us. Perceptually: one-meter experiences represent things as having that property that is the normal cause of one-meter-ish experiences in us.

As before, size functionalism is not the only view of the case. One could be a size presentationalist, one which all one-meter-ish experiences represent a size of one meter. As before, there will be internalist realist presentationalism (on which at most one of Oscar and Doubled Oscar is perceiving veridically), externalist realist presentationalism (on which both may be perceiving veridically because they are having different experiences), and eliminativist presentationalism (in which neither is perceiving veridically because the represented size properties are never instantiated). I do not think any of these views are very attractive. As before, internalist realism tends to lead to size skepticism, and it is not easy to see how the relevant relation to sizes can be established. Externalist realism suffers from the usual problems of indeterminacy, and it is also not easy to see what the relevant difference in size experience would be like phenomenologically. The view that all size experiences represent uninstantiated properties seems unattractive and unmotivated by the phenomenology.

As in the case of orientation, one can define p-categoricalism about size: the view that there are categorical size properties in physics. One can cash this out as the view that there are distinct possible worlds physically isomorphic to this one except that distances there are twice as large. P-categoricalism about size is widely rejected in contemporary physics: a putative universe just like ours but twice the size would in fact be in the same physical state as our universe. Instead, physics

suggests p-relationism about size: there is merely relative size in physics, not absolute size. To be sure, there may be constants such as the speed of light tied to size, but these play the same role as parity violations for orientation. They do not establish categorical sizes, since a universe twice the size (with speed of light twice as fast) would still be a copy of our universe.

One can also define e-categoricalism about size: the view that there are categorical experiences of absolute size. One can cash this out as the view that there are total experiences of size phenomenally isomorphic to normal size experiences but in which the phenomenology of size is doubled: one-meter-ish experiences are replaced by two-meter-ish experiences, and so on. One can get a limited sort of phenomenological doubling from the two different body-doubling scenarios considered earlier, where one case one's body seems to double in size and the environment stays constant, and in the other case one's body seems to stay constant and the environment halves. It is arguable, though, that this difference arises from holding fixed the representation of some sizes, such as the previous size of one's body. In one scenario, my new body is represented as the same size, while in the other it is represented as twice the size. E-categoricalism requires that all of our size representation (including memory representation and the like) could double in this way. It is at least far from clear that this is possible. On the face of it, our phenomenology does not seem to reveal absolute sizes in the way it seems to reveal absolute colors. Instead it seems only to directly reveal relative sizes.

As in the case of orientation, if e-categoricalism about size is false, then all forms of size presentationalism are false. If p-categoricalism about size is false, then all realist forms of size presentationalism are false. If e-categoricalism is true but p-categoricalism is false, the door is opened for an eliminativist form of presentationalism, but the view that all size experience is illusory is not attractive. So I think there is good reason to reject size presentationalism and accept size functionalism.

A natural fallback is to accept that absolute size properties are not presented in experience, while holding that *relative* size properties are presented in experience. I discuss this in the third puzzle concerning shape.

Before moving on: are doubling scenarios physically possible? A first approximation to an answer is that they are physically possible in classical physics but not in contemporary physics. In classical mechanics with point particles, if a certain scenario is physically possible, an isomorphic scenario involving particles of the same masses with doubled distances between the particles and doubled velocities will also be physically possible. The doubled scenario will evolve in an isomorphic way over time. The same goes for classical mechanics augmented by a single force,

as in the Newtonian theory of classical mechanics plus gravitation. Something similar applies to classical continuum mechanics. We simply double all the lengths and velocities, reduce densities by a factor of 8 to compensate, and the resulting system will evolve in an isomorphic way.

To a second approximation, there is more to say about classical physics.⁵ It is true that any physically possible scenario can be doubled. But it is not true that there are physically possible doubling scenarios involving human bodies and the like. The reason is that scenarios involving human bodies and the like are not really physically possible in classical physics. Classical physics does not have the resources to explain the cohesion of matter: why bodies stay together rather than having their constituents fly apart. Some classical physicists postulated extra spring force laws to explain the cohesion of matter, but these laws then rule out doubling scenarios. A correct explanation of the cohesion of matter was only given by quantum mechanics, in which doubling scenarios are also not physically possible.

One might also think that biological scaling principles block doubling scenarios. It is well-known that large animals are subject to constraints quite different from those of small animals, for example, and their dynamics are far from isomorphic. It is plausible, however, that the key differences here are explained by factors that remain constant: the gravitational field, the density of water, the density of matter, and so on. In a Doubled Earth scenario, these factors will change in a way that they do not in actual biology, so the facts about biological scaling alone do not block these scenarios. That said, facts about relativity (the fixed speed of light, say) and quantum mechanics do seem to block the scenarios.

How can we use arguments involving doubling scenarios if they are physically impossible? The most obvious way is to appeal to *approximate* doubling scenarios that are physically possible. For example, even if there cannot be a duplicate of a human who is twice the size, there can still be humans who are quite different in size from each other, with relevantly similar brains and relevantly similar experiences. One can pose similar questions about these humans: is one undergoing a size illusion, or another? Our verdicts here will plausibly be analogous to our verdicts concerning Doubled Earth: for example, that neither is undergoing size illusions. From here can can draw conclusions similar to those we draw from Doubled Earth: for example, that the similar experiences do not present absolute sizes. An opponent can always suggest that the differences between the cases make a relevant difference: for example, perhaps the similar experiences differ in a way that enables them to present different absolute sizes. But it will typically be implausible

⁵Thanks to Mark Wilson for discussion here.

that these differences are the sort of thing that would explain the differences they need to explain.

It is also worth noting that the appeal to doubling scenarios is really a concession to the opponent of spatial functionalism, who wants to argue that it is epistemically possible that things are twice as large (or just as large) as they seem. This opponent is committed to the epistemic possibility of certain doubling scenarios. The mere claim about epistemic possibility is consistent with the physical impossibility of these scenarios, as the relevant laws of nature are not themselves *a priori*. And we can at least play along with those claims, considering what will be the right thing to say if it turns out that our laws permit doubling scenarios and that we in fact inhabit such a scenario. The correct answer, at least in worlds reasonably close to our own, seems to be that in those scenarios we do not suffer from illusions. This is enough to suggest that size presentationism is false, and that in these scenarios, size concepts picks out whatever normally causes the right sort of size experience. The epistemic profile of size concepts does not depend on questions of physical possibility, so conclusions about these epistemic profile are robust on whether or not doubling scenarios are physically possible.

One might object that my arguments have appealed to empirical premises about actual physics, and that I cannot simultaneously use actual physics and ignore it. However, I have used actual physics only for the purpose of arguing that *p*-categoricalism is false. We can then use relatively *a priori* arguments to consider what follows from the falsity of *p*-categoricalism. For example, we can first argue that in worlds where *p*-categoricalism is false but local doubling is possible, doubled scenarios will not involve illusions, so there are no permanent doubling illusions. In worlds where *p*-categoricalism is false and local doubling is impossible, there can also be no permanent doubling illusions. So if *p*-categoricalism is false, there are no permanent doubling illusions.

As in the case of orientation, my arguments for *p*-categoricalism do not establish it with certainty, so the ensuing arguments do not rule out permanent doubling illusions with certainty. For all those arguments establish, it remains epistemically possible with we live in a universe with absolute size (perhaps with misleading evidence against absolute size) and in which there is a doubling illusion. However, I think there is a very strong case against *e*-categoricalism about size, which (even more than in the case of orientation) may then allow us to rule our permanent size illusions with certainty. So we have good and perhaps conclusive reason to accept size functionalism and reject size categoricalism, and conclusive reason to accept that if size functionalism is true, there are no permanent size illusions.

Third Puzzle: Shape Illusions

The third puzzle asks: is it possible that you have been under a lifelong shape illusion where everything is twice as wide (in a certain direction) as it seems to be: things that appear to be square are actually 2:1 rectangles, things appear to be spherical are actually elongated ellipsoids, and so on?

As in the first puzzle, the sort of possibility that is most relevant is epistemic possibility, and the sort of illusion that is most relevant is perceptual illusion. Strictly speaking an illusion involving a fixed direction in space may make the most sense here, but for purposes of illustration we can suppose the illusion involves a fixed direction relative to the Earth, at least for a local environment. It is also worth noting that this shape illusion comes along with an illusion about relative size: when a square is misperceived as a 2:1 rectangle, one side is perceived as twice as long as another when in fact they have the same length.

As before, it is certainly possible that I am under a temporary size illusion of this sort. We could devise compressing lenses that compress the image horizontally in my visual field, for example. At least when I start wearing such lenses, I will be under a visual illusion along the lines above: squares will look like 1:2 rectangles, and so on.

What about permanent shape illusions? What if we had fitted a baby since birth with these compressing lenses, along with analogous devices for other sensory inputs and for motor effectors. Would the resulting person suffer from life-long illusions? More generally, can we make sense of the hypothesis that we are undergoing illusions of this sort? Intuition is less clear here than for the first two puzzles: many more people are at least tempted to say yes.

We can clarify a case as before by appealing by appealing to a Twin-Earth-style thought experiment. Here the relevant case is El Greco world, introduced by Susan Hurley (1998) and discussed in this context by Brad Thompson (2010). On El Greco world, everything is just like our world, except that things are elongated by a factor of 2:1 in a certain fixed direction. The laws of dynamics are changed in a corresponding way so that the way things develop on El Greco World is isomorphic to the way it develops on our world.

People sometimes question whether El Greco World with isomorphic dynamics is really coherent. One way to see that it is coherent is to note that dynamics on El Greco World will look just like the dynamics of a stretched-out movie in our world. Bodies that are rigid in our world will correspond to nonrigid bodies on El Greco world: for example, as a square turns 90 degrees clockwise on Earth, its counterpart on El Greco world will start as a 2:1 rectangle, then its sides

will gradually equalize until it is square at the 45 degree mark, then shortening and lengthening will continue until the side that was initially twice as short as the other is now twice as long as the other. Another way to see that it is coherent is to note that this sort of relative compression and elongation of two scenarios is actually physically possible under the theory of special relativity, as I will discuss shortly.

On Earth, Max sees a square. On El Greco World, Max's counterpart Twin Max is an elongated functional duplicate of Max. He sees (what we would call) a 2:1 rectangle, and says 'That's square'. Is Twin Max mistaken? Is one of Max and Twin Max undergoing a shape illusion? Both? Neither?

As before, intuitions about the case of shape are less clear than intuitions about the cases of orientation and size. In this case there is a stronger temptation to say that Twin Max is undergoing an illusion: he sees something as square when it is actually far from square. On this view, Max and Twin Max are having the same sort of squarish experience, and all squarish experiences represent squareness. Only Max sees a square, so Max is perceiving veridically and Twin Max suffers from an illusion.

Thompson (2010) argues that this view is mistaken, and I agree. I think one can support Thompson's verdict with an appeal to a physically possible case deriving from the special theory of relativity.

The case uses the well known phenomenon of Lorentz contractions. According to the special theory of relativity, if an object is at about 0.87 times the speed of light relative to us, its length will contract by a factor of 2 in the direction of travel. For example, a stick that is one meter long in the direction of travel will contract to a half-meter long. A sphere will contract to become a compressed spheroid. And so on.

We can now imagine Lorentz Earth, which is a physical duplicate of Earth, except that it is traveling at 0.87 the speed of light relative to Earth—say, on the plane of its equator. Then according to special relativity, where Earth is roughly spherical, Lorentz Earth is compressed so that it is roughly ellipsoidal. People on Lorentz Earth are functional duplicates of people on Earth, compressed 2:1 in the direction of travel. Around the poles, people will be unusually thin by Earth standards, while around the equator, they will be unusually short.

Now, suppose that on Earth, Albert is at the North Pole and sees (what we call) a square. On Lorentz Earth, his counterpart Twin Albert sees what he calls a 'square' but what we call a 2:1 rectangle. Is one of Albert and Twin Albert suffering an illusion? Both? Neither?

In this case, it is extremely implausible that exactly one of Albert and Twin Albert is having

an illusory experience. From an objective point of view, the situation is completely symmetrical. Recall that in special relativity there is no absolute reference frame. Rather, Twin Albert is compressed relative to Albert's reference frame, while Albert is compressed relative to Twin Albert's reference frame. It is plausible that whether someone is undergoing an illusion is not itself dependent on a choice of reference frame. If so, then if Albert is having a veridical experience, so is Twin Albert, and vice versa. It is independently plausible that Albert is having a veridical experience. So the natural verdict is that both are having veridical experiences.

This is already enough to suggest that 'square' is Twin-Earthable. Albert uses 'square' to pick out (what we call) squares. Twin Albert uses 'square' to pick out (what we call) 2:1 rectangles. The rectangle that Twin Albert is seeing is in the extension of his word 'square', but it is not in the extension of Albert's word 'square'.

One might suggest in response that there is something non-Twin-Earthable nearby. In particular, we can define 'rest square' so that an object is rest-square (at a time) iff it is square relative to a frame of reference in which it is at rest (at that time). Then the objects seen by Albert and Twin Albert are both rest-square, and nothing about this scenario suggests that the word 'rest-square' is Twin-Earthable.

A presentationalist might then suggest that 'square' refers to rest-squareness for both Albert and Twin Albert, and that squarish experiences represents rest-squareness for both. But at least as English is ordinarily used, 'square' does not seem to mean rest-square, just as 'length' does not seem to mean rest-length. This is witnessed by the fact that Lorentz contractions are usually described as Lorentz *contractions*: sticks do not stay the same length as they accelerate, and squares do not remain square. If these claims are correct, the squareness they refer to is not rest-squareness. Similarly, squarish experiences do not seem to represent rest-squarishness. If they did, we would have to say that the rectangularish experience we have when observing an accelerated rest-square object is illusory, but this verdict seems contrary to Einstein's standard analysis of the perception of moving objects.

To say that 'square' is Twin-Earthable is not yet to say what property it refers to. For example, 'square' (and square experiences) might pick out squareness relative to one's current reference frame, or relative to one's historical reference frame, or something else. The historical-frame view seems implausible, since it suggests that someone who accelerates to Lorentz Earth and has squarish experiences of the rest-squares there will be undergoing an illusion. The current-frame view delivers a more plausible result in this case, and I think it is the most plausible view overall. There is a tricky issue about whether the current-frame view must say that one's memories of Earth

in such a case are mistaken; perhaps one can give special treatment to the content of memories. In any case, for our purposes it is enough to establish that shape concepts are Twin-Earthable and that shape presentationalism is false.

To take things a step further, we can introduce the case of Absolute Lorentz Earth. This requires a counterfactual physics, which is just like special relativity except that there is an absolute reference frame that defines absolute rest. Most physicists think there is little reason to think this theory is true, but it is certainly coherent and consistent.

We can now run the thought-experiment above using Absolute Lorentz Earth. Albert sees what he calls a square, while Twin Albert sees what he calls a square. For a twist, let us suppose that Absolute Lorentz Earth is actually at rest in the absolute reference frame, while Earth is traveling at 0.87 times the speed of light relative to that frame. So Twin Albert is seeing an object which is square relative to the absolute frame, while Albert is seeing a 2:1 rectangle relative to that frame. Should we say that Albert is undergoing an illusion?

My own judgment is reasonably clear: Albert is not undergoing an illusion. We have already seen that in the Lorentz Earth case, Albert's experience is veridical. The Absolute Lorentz Earth case is relevantly like that case, and the introduction of an absolute frame does not seem to be the sort of thing that could change Albert's experience from veridical to illusory.

To back up this judgment, we can consider what we would say if we discovered that we were living in a world traveling 0.87 times the speed of light relative to the absolute reference frame. Would we say that all our size experiences are illusory and that our statements such as 'That is square' are false? Perhaps a few philosophers might take that line, but I predict that it would be widely rejected by the community. Instead, we would probably distinguish (say) absolute squareness from ordinary squareness (squareness relative to our reference frame), and we would say that our term 'square' refers to ordinary squareness. Likewise, we would say that the veridicality of our squarish experiences turns on whether the object is square in the ordinary sense, not in the absolute sense. I think that this verdict would be the right verdict.

Once we have this verdict about Absolute Lorentz Earth case, it is not hard to justify a similar verdict about El Greco world. The two cases are relevantly similar. If we turn out to be in a world that is elongated relative to the rest of the galaxy, we might distinguish (say) ordinary squareness from galactic squareness, but we would say that our term 'square' refers to ordinary squareness. Even if it turns out that our squares are elongated according to a measure in fundamental physics, we would distinguish (say) macrosquareness from microsquares, and insist that our term 'square' refers to macrosquareness. I think that this verdict would be correct, and

that a corresponding verdict about the veridicality of our squarish experiences would be correct.

If all this is right, then 'square' is clearly Twin Earthable. Perhaps we can define an expression such as 'local square' (analogous to 'rest square') that is not, but that will just be something along the lines of 'square by local standards', which comes to 'the sort of thing that causes square experiences in its vicinity'. (Compare: 'locally red' = 'the sort of thing that causes reddish experiences in its vicinity'.) But importantly, the ordinary word 'square' will be Twin-Earthable. Equally importantly, perceptual experiences will be equally veridical on Earth and in all the twin-earth scenarios.

All this tends to support *shape functionalism*: shapes such as squareness are picked out in virtue of their role in causing our experiences of shape. In fact, these cases suggest that our relevant concept of shape is a concept of whatever normally causes the relevant shape experiences.

As before, shape functionalism is not the only possible view of the case. One could also be a shape presentationist, on which all squarish experiences represent the same property: squareness. The internalist realist presentationist will hold that at most one of the twins perceives veridically: presumably Twin Albert in the Absolute Lorentz Earth case, and Max in the El Greco case. But it is hard for this view to give a clear verdict about the Lorentz Earth case. Perhaps it will hold that squarish experiences represent rest squareness here (so both Albert and Twin Albert perceive veridically), but this is then hard to reconcile with standard relativistic intuitions about the perception of fast-moving objects, as well as with verdicts in the Absolute and El Greco cases.

The externalist realist presentationist can hold that Max and Twin Max are having different sorts of experience: Max is having a squarish experience while Twin Max is having a rectangle-ish experience, and both are veridical. Again, it is hard for this view to deliver a coherent verdict about Lorentz Earth that squares with these other verdict, though. Another oddity is that Twin Max's rectangle-ish experiences seem to be mismatched with his behavioral patterns of action and judgment: these patterns have a symmetry that suggests a square-perceiver, and the rectangularity of his experience does not seem to be reflected in his behavior.

Finally, the eliminativist presentationist will hold that squarish experiences represent Edenic squareness, a property that nothing has in the actual relativistic world. So both Albert and Twin Albert are perceiving veridically in the Lorentz Earth case. Perhaps there is room for this view to allow that one of them gets it right in the Absolute Lorentz Earth case and in the El Greco case (assuming this arises within a Newtonian world). But in the actual world, all squarish experiences are illusory. I think that this view has some attractions, just as the corresponding view has in the case of color experience, and much more than the correspond views of size and orientation.

Nevertheless, it is far out of kilter with our actual judgments about shapes and shape experience in our relativistic world. It is more plausible to hold that ‘square’ refers not to Edenic squareness but to ordinary squareness, the property that plays the associated role in the actual world. We could still say that for a squarish experience to be perfectly veridical, an Edenic square would be required, but ordinary (imperfect) standards of veridicality require only ordinary squareness.

Presentationalist intuitions get a little more grip for shape than for orientation and size, reflecting the fact that e-categoricism about shape is much more plausible than about size and orientation. There does seem to be a categorical difference between a squarish experience and a rectangle-ish experience, just as there is between a reddish experience and a greenish experience. As for p-categoricism about shape, this would be true in an Newtonian world. In our relativistic world things are more complicated: shape in the ordinary sense is relative, there are still rest shape properties that are categorical, but neither seems a good match for what is directly presented in our categorical shape experience. Instead we are left in a situation analogous to the color case: shape experiences directly present Edenic squareness, which is not instantiated in our world, and they represent ordinary squareness, which is instantiated in our world.

2 Spatiotemporal functionalism

All of this tends to suggest a thoroughgoing spatial functionalism, at least about our relation to the spatial properties instantiated in our world. The size, shape, and orientation properties instantiated in our world are not directly presented in our shape experience. Instead, we pick out these properties in virtue of the causal role they play in bringing about our spatial experience.

Of course there is more to space than size, shape, and orientation. There are other spatial features: betweenness, continuity, dimensionality, for example. Although it would take more work to show this, I think that an analogous moral applies in all these cases too. In each case there is a possible Twin Earth case in which a corresponding term refers to something else. (One way to construct such cases comes from the reflections in “The Matrix as Metaphysics” on Matrix cases and pre-spatial computational levels in physics.)

One can also construct analogous cases involving time. One such case involves time-inverted societies who are as near to the Big Crunch as we are to the Big Bang: it is plausible that their term ‘past’ refers to the same relation that we call ‘future’. There are also speeded up subjects for whom ‘one second’ refers to a half-second. All this suggests that temporal expressions are also Twin-Earthable, and that the corresponding properties are not directly presented in tempo-

ral experience. As in the case of shape, there does seem to be a categorical difference between past-ish and future-ish experiences, and correspondingly Edenic pastness and Edenic futureness may be directly presented in these experiences. But there is not much reason to think that these relations are instantiated in the actual world or that they are required for the ordinary veridicality of our temporal experience. So we are left with a sort of temporal functionalism, on which instantiated temporal properties and relations are picked out in virtue of their relations to our temporal experiences.⁶

As in the case of color, we have in effect moved from primitivism to functionalism. At one point it might have been reasonable to think that our world contained primitive Edenic color properties, but science suggested that it does not. As a result we have come to accept color functionalism rather than color primitivism about the color properties instantiated in our world. Likewise, at one point it might have been reasonable to think that our world contained primitive Edenic spatial or temporal properties, but science suggested that it does not. As a result, we should accept spatial and temporal functionalism rather than spatial and temporal primitivism about the spatiotemporal properties instantiated in our world.

Spatial functionalism helps in another domain: helping understand the role of space in contemporary physical theories in which ordinary space is not fundamental. On a fairly standard view of quantum mechanics, for example, the fundamental entity is a wave function, and the fundamental space that it inhabits is not three or four-dimensional space but a high-dimensional configuration space. Ordinary spatial properties must then be somehow derivative from this configuration space, and the question arises as to how. A natural view is that they are picked out in virtue of their functional role, and in particular they are picked out as those nonfundamental properties that serve as the causal basis for our spatial experience. The same applies to other physical theories in which ordinary space is not fundamental, such as string theory.

⁶If even spatial and temporal expressions are Twin Earthable, one might wonder which expressions are not. This is too large a topic to settle here, but in *Constructing the World* I suggest the following list: logical and mathematical expressions, mental expressions such as 'conscious', nomic expressions such as 'law of nature' and possibly terms tied to causation and counterfactual dependence, and some structural expressions such as 'object', 'property', and 'fundamental'. There are also expressions for Edenic properties, such as 'Edenic redness' and 'Edenic squareness': these are not Twin-Earthable, but the properties they pick out are not instantiated in our world.

3 Anti-Skeptical Conclusion

If spatial primitivism were correct, it would not be hard to sustain systematic lifelong illusions about space. We would attribute a certain distribution of primitive properties to the world, and could easily make sense of hypotheses on which this distribution is not as we take it to be. In fact, it is arguable that the actual non-Newtonian world would be one in which by the lights of spatial primitivism, spatial experiences would be massively nonveridical.

If spatial functionalism is correct, on the other hand, then systematic lifelong illusions about space are much more difficult to sustain. In particular, if we pick out spatial properties as the normal causes of spatial experiences, then situations in which spatial experiences are normally caused by properties other than the spatial properties they represent will be ruled out. There may still be room for some permanent illusions if this relation works in a holistic way: if we fix reference to a manifold of spatial properties all at once, as the normal causal basis of a manifold of spatial experiences, then there is room for the occasional normal mismatch between an experience and the property it represents. There may also be cases in which there is no normal cause of a certain sort of experience. Still, it will at least be much harder to get permanent illusions off the ground.

All this has consequences for standard ideas about skepticism. The intuition that Cartesian skeptical scenarios—evil genius vases, Matrix cases, and so on—involve systematic deception turns largely on the intuition that they involve spatial illusions. If a subject has veridical spatial experiences and true spatial beliefs, they can hardly be said to be systematically deceived about the external world. So if lifelong spatial error is impossible, these lifelong skeptical scenarios is impossible too. And if lifelong spatial error is hard to sustain, lifelong skeptical scenarios are hard to sustain too.

In particular, spatial functionalism suggests a view on which subjects in these putative skeptical scenarios in fact have veridical spatial experiences and true spatial beliefs. Their spatial experiences will pick out the normal causes of these experiences, and there will certainly be normal causes in these scenarios: in the Matrix case, for example, these will be computational properties in the Matrix. The same goes for spatial expressions. If so, then the Matrix scenario will not be a permanent error scenario at all.

Of course the matter is not completely cut and dried. One could argue that the Matrix case is more distant than the spatial Twin Earth cases discussed earlier, and the same morals do not apply. For example, perhaps the Matrix distorts underlying structure involving continuity, betweenness,

dimensionality, or some other feature. Or perhaps spatial functionalism requires a relation more complicated than the “normal cause” relation to pick out truly spatial properties. I think the gap here can be bridged by the kind of reasoning in “The Matrix as Metaphysics”. But the analysis of spatial functionalism given here helps us to see what is really going on in these cases.

This is not to say that the skeptic is entirely refuted. Spatial functionalism allows temporary error scenarios for any specific spatial experience or belief, and temporary error scenarios will be enough for some skeptical purposes. There will also be some lifelong error scenarios in which there is no normal cause: one in which random input causes regular experience completely by accident, for example. There is also the point discussed earlier that the argument for against spatial primitivism and for spatial functionalism is partly empirical: so we cannot rule out skeptical scenarios in which we are living in an Edenic world with misleading evidence that we are in a non-Edenic world, and in which we are globally deceived in other ways. Still, it seems that spatial functionalism at least softens up previously strong skeptical intuitions, and leaves them less robust before.

I conclude, somewhat paradoxically: precisely because we are directly presented with fewer features of the world than we might have thought, we are less open to illusion and deception.

Appendix: What are Left and Right?⁷

What sort of entities are left and right? We use these terms in different ways: we say that a hand is a left hand, that cars drive on the right (or left) side of the road, that on the Australian flag the Union Jack is to the left of the Southern Cross. To make matters more complicated, left and right seem to be relative to a perspective: the right side of the road for one driver is the left side for a driver coming in the other direction, the Union Jack is to the right of the Southern Cross from the perspective of someone behind the flag, and even the leftness of a left hand seems to reflect the perspective of the subject whose hand it is.

It is natural here to invoke a three-place relation $L(o1, o2, o3)$ between two objects and an observer, holding when object $o1$ is to the left of object $o2$ relative to perspective of $o3$. We can then say that a left hand is a hand belonging to an arm that is to the left of the associated body from the perspective of the associated subject. The left side of the road is the side that is to the left of the center line from the perspective of a driver. The Union Jack is to the right of the Southern Cross from the perspective of someone viewing the flag in its standard arrangement.

⁷This appendix doesn't appear in the published version.

For our purposes we can simplify by focusing on the case where the object o_2 and the perspective o_3 are identical. This yields a two-place relation between an object o_1 and an observer o_2 , holding when o_1 is to the left of o_2 from the perspective of o_2 . It is not obvious whether one can define the three-place relation in terms of the two-place relation, but there are some prospects for doing so using the tools below, and in any case the two-place relation is enough to raise all the issues in this paper.

Underlying this two-place relation seems to be a two-place relation between a point and a perspective, holding when the point is on the left, relative to a perspective. What is a perspective? It plausibly involves a point, but it is not just a point: the same side of the road can be to the left or to the right of the center line from the perspective of someone at a given point, depending on which way they are facing. To fix a perspective, it seems that we have to fix not just a point, but also which direction is front and which direction is up. With those fixed, matters of left and right are fixed. We could regard a perspective as a point along with two rays from the point representing front and up, or simply as an ordered triple of three points.

We could then understand leftness as fundamentally involving a relationship $L(p_1, p_2, p_3, p_4)$ between four points, holding when p_1 is to the left of p_2 from the perspective of p_2 when p_3 is taken as front and p_4 is taken as up. Or perhaps more elegantly, we can understand it as a relation $L(r_1, r_2, r_3)$ between three rays emanating from a single point, holding when r_1 is to the left from a perspective where r_2 is taken as up and r_3 is taken as front.

Appendix: Is There A Categorical Phenomenology of Left and Right?

In the discussion of the first puzzle, I defined p-categoricalism (physical orientation categoricalism) as the thesis that there is a categorical difference between left and right in the physical world. More precisely, p-categoricalism is the thesis that there is a possible world distinct from this one that is a physical duplicate of this world except that it is left/right reversed. (So there might be distinct worlds, one containing only a left-handed glove and one containing only a right-handed glove.) P-relationism is the denial of p-categoricalism, and in effect is the view that there is merely a relational difference between physical left and right.

Likewise, I defined e-categoricalism (experiential orientation categoricalism) as the thesis that there is an categorical difference between the experiences of left and of right. More precisely, e-categoricalism is the thesis that for normal human total experiences, there is a possible experience that is a phenomenal duplicate of it except that it is left/right reversed. [Q: Should one also add

something about having left/right inverted memories/associations?.] E-relationism is the denial of e-categoricism, and in effect is the view that there is merely a relational difference between leftishness (the experience of left) and rightishness (the experience of right).

E-relationism tends to suggest that our grip on the difference between leftishness and rightishness is fundamentally demonstrative (or indexical): leftishness is *this* orientation, and rightishness is *that* phenomenal orientation. It also tends to suggest (especially in conjunction with p-relationism) that our grip on left and right is demonstrative in a similar way (either directly—left is *this* orientation, or indirectly, left is what causes *this* phenomenal orientation). By contrast, e-categoricism tends to suggest that we have a substantial non-demonstrative grip on the difference, akin to the grip we plausibly have on different color experiences, while p-categoricism suggests something similar about our grip on external relations.

In the text, I offered considerations from physics in favor of p-relationism. Here I will offer some considerations in favor of e-relationism. These considerations are partly grounded in cases Geoffrey Lee discusses in “The Experience of Left and Right”, in which he endorses a kind of e-categoricism, and argues from there to p-categoricism.

Lee does not argue at length for e-categoricism in his paper. He largely takes it as obviously correct and uses it to argue for other conclusions. He gives one simple argument for the view: if I were presented with a mirror-reversed environment, I would have experiences that are left-right reversed with respect to mine.⁸ But a moment’s reflection suggests that this is not obvious. Perhaps there will be inversion at the level of sensory qualities, but there may also be many differences. Text going from right to left will be unreadable for me, for example, and this will yield a phenomenal difference. Cars will seem to be on the wrong side of the road and familiar objects will seem the wrong way around. Even without these familiar objects, bodily asymmetry will yield experiential asymmetry: I am right-handed, so if I play tennis with my left hand, the experience will be entirely different.

Lee acknowledges these differences but suggests that that the difference between leftishness and rightishness must go deeper than these associations. One way to test the point is to imagine inverting those associations too: perhaps my body is flipped too, I read from right to left, my memories are flipped. On Lee’s view, the resulting subject will then be left-right inverted with respect to me. I think that this is no longer obvious, though. Perhaps what it is like to be that subject is exactly the same as what it is like to be me.

In fact, Lee gives a separate argument that can be used to argue directly from p-relationism to the falsity of e-categoricism. He considers a subject, Simon, with a symmetrical brain in a

symmetrical universe. Simon's experiences will be initially symmetrical. Then an asymmetry is introduced: perhaps a quantum fluctuation produces a flash in the environment on one side of him, and he experiences a flash on one side of his visual field. Although it might seem that the flash could be introduced on one side or the other, p-relationism says that there is just one possible resulting physical state. Likewise, under plausible assumptions, e-relationism says that there is just one possible resulting experience, while e-categoricism seems to require that there are two possible resulting experiences. The combination of e-categoricism with p-relationism therefore at least threatens to violate the supervenience of phenomenal states on physical states.

One can argue that the following four theses are inconsistent: (1) e-categoricism (experiential left/right categoricism), (2) p-relationism (physical left/right relationism), (3) physicalism (global supervenience of phenomenal properties on physical properties), and (4) orientation dependence (phenomenal orientation depends on physical orientation, so that physical reversals yield phenomenal reversals). If p-relationism is correct, then the physical states of the world in the left-flash and right-flash scenarios are identical. If e-categoricism and orientation dependence are correct, then the phenomenal states of the world in these scenarios are distinct. These claims are inconsistent with physicalism.⁹

Lee uses this structure to mount a "transcendental argument" from e-categoricism to p-categoricism. But we have seen that p-categoricism is widely rejected by physicists and philosophers of physics, and that physical relationism is widely regarded as much more plausible. So one can equally use this structure to argue from p-relationism to e-relationism: there is just one possible physical state in this situation, and one possible phenomenal state.

Of course this argument could be resisted in various ways. One could embrace p-categoricism, as Lee does; but this requires denying a widely accepted view of physics. Lee invokes parity violations to support his p-categoricism here, but as Baker (2011) notes, these violations do not really support p-categoricism over p-relationism.

One could also deny physicalism, holding that phenomenal states here do not globally supervene on physical states. I deny physicalism myself, but this does not suffice to block the argument. One can run a version of the argument that merely appeals to global *nomological* supervenience on the phenomenal on the physical, which I accept. At least if physics were symmetrical (without parity violations and the like), then the Simon scenario would be nomologically possible. Then while p-relationism suggests that the left-flash and right-flash scenarios involve just one

⁹Thanks to Adam Pautz for discussion here.

physical state, e-categoricism plus orientation dependence suggest that they involve two distinct phenomenal states. This violates global nomological supervenience. So in a world without parity violations, global nomological supervenience plus p-relationism plus orientation dependence entails e-relationism. In a world with parity violations one will not get direct counterexamples to nomological supervenience, as mirror-reversed physical scenarios will not be nomologically possible. Still, the main loopholes opened for the p-relationist who is an e-categoricist are to embrace psychophysical laws that connect categorical phenomenal orientation to weak interactions or to non-orientational physical features (denying orientation dependence), neither of which seems an especially attractive view. So even global nomological supervenience tends to support the argument.

Finally, one could hold that e-relationism is true of Simon but that e-categoricism is true of us.¹⁰ Lee entertains the idea that initially symmetric beings such as Simon will have fundamentally different experiences from those that we have, in that they will not have categorical orientations. Still, it seems plausible that the right course of asymmetrical development could evolve to give Simon an asymmetrical brain that is very much like that of a normal human, with asymmetrical experiences that are at least isomorphic to those of a normal human. If e-categoricism were true of Simon's experience even at this late date, then consideration of a left/right reversed developmental history would still lead to inconsistency with p-relationism and supervenience. So the proponent of this line must hold that e-relationism is still true of Simon at this late date, and that his experiences differ from our isomorphic experiences in lacking categorical orientation. I do not have any clear grip on what this difference would consist in, and think it is much more plausible that Simon's resulting experience would be just like our own. If so, then if e-relationism is true of Simon, it is true of us.

Lee does not seriously entertain the possibility that we are e-relational. Instead, he takes related ideas as *reductio* to be avoided. In his discussion of the related but distinct notion of "experientially symmetrical" beings (beings that have identical experiences of mirror-reversed environments, like Simon in his initial state), he suggests that these beings would be unlike us in

¹⁰I think that moves like this are at least sometimes correct. For example, I think an analog of this claim is the most plausible verdict in Tim Schoettle's (2009) extremely interesting color-inversion scenario involving someone (Perry) with a color-processing architecture that is quite different from ours. I think that most plausible view of Schoettle's case is that Perry does not have categorical color experiences, as we do, but merely has relational color experiences, experiencing certain surfaces as the same or as different in various respects. But Perry's cognitive architecture is vastly different from ours, while Simon's cognitive architecture may be just like ours.

that they would not experience the world as oriented. I am not sure this is correct. Even Simon can experience the world as having two orientations. There is just no categorical difference between the two orientations. Over time, non-categorical differences between the two orientations will build up due to asymmetrical associations, dispositions, and the like. After these associations have built up, Simon will no longer be experientially symmetrical: a reversed environment will trigger different experiences due to different associations. But he will still have the related property of being e-relational, in that his total experience would be the same if he were in a mirror-reversed environment with mirror-reversed associations, memories, and so on.

I do not see any clear reason to deny that we are e-relational in this sense. Of course being e-relational is compatible with all sorts of asymmetrical perception, behavior, associations, and memories. The situation is akin to that of p-relationality in physics, which is perfectly compatible with asymmetrical dynamics in the fundamental laws and all sorts of asymmetrical phenomena at the macroscopic level.

Why reject the hypothesis that we are e-relational? One reason might be a brute phenomenological intuition of categoricity: but I think this intuition is at best extremely weak (certainly much weaker than in the case of color) and not hard to deny. Another reason stems from the conceivability of experiential inversions, such as mirror-reversed environments: but these merely invert one aspect of experience, as we have seen, and it is much less clear that total experience could be inverted. Another reason stems from the intuition that the differences between leftishness and rightishness in our experience cannot stem just from associations, memories, dispositions, and the like; but I do not see much reason to accept this.

One reason to take seriously involves successive experiences. Say that Simon starts in his symmetrical state, then experiences a flash on one side, then on the other. Surely these two experiences will be distinct, not identical. The e-relationist can appeal to differences in memory: the second flash will be experienced as on a different side from the earlier flash. What about a case without memory? The intuition of phenomenal difference is now not as strong, but to respect the intuition, the e-relationist can appeal to the continuity of a stream of experience over time. Simon has a temporally extended stream of experience involving a flash on one side and then on the other side. This observation is quite compatible with e-relationism, as long as we allow that he would have had a phenomenally identical stream of experience had the flashes come in the reverse order. Arguably this saves the intuition of a difference in experience, though the difference is now revealed as holistic: the experiences are locally phenomenally identical, but differ from each other in virtue of the way they are embedded in a stream of consciousness. The situation is analogous to one

involving simultaneous symmetric experiences: these experiences are locally phenomenally identical, but differ from each other in virtue of the way they are embedded in a (synchronic) total state of consciousness. In effect, we can understand the total states in the definition of e-relationalism to be full streams of consciousness, thereby allowing local differences between mirror-reversed experiences within a stream of consciousness.

What if we move to a case with no memory and no stream of consciousness? Simon experiences a flash on one side, then forgets the flash (perhaps becoming symmetrical again) and becomes unconscious for a moment, then experiences a flash on the other side. In this case, I think the intuition that there is a phenomenological difference between the two experiences is not too hard to reject. At least it seems no harder than in the case of counterfactual experiences, where there is a flash on one side but there could have been a flash on the other side. The plausible thesis of p-relationism entails that the original situation and the putative counterfactual situation are in fact physically identical. This thesis is perhaps a little counterintuitive, but once one has accepted it, it is not a big stretch to accept that the original experience and the putative counterfactual experience are in fact phenomenally identical. Having gotten this far, it is not hard to accept that the two successive experiences (without memory or continuity) are also phenomenally identical.

If one wants to hold on to the intuition of phenomenal difference in this case, one could also embrace a variant of e-relationism on which there can be phenomenal differences between mirror-reversed total experiences, but only within a single subject. In effect, there is a way of aligning phenomenal orientations over time within a subject, but not across subjects. This leads to a version of the Frege-Schlick view (which Lee acknowledges): across different subjects, there may be no fact of the matter whether two experiences are phenomenally the same, or mirror-reversed. On balance I think I prefer to reject the intuition of phenomenal difference in this case, but the Frege-Schlick view also deserves attention.

One can distinguish the views by considering near-symmetrical beings A and B, both of whom experience a red flash on one side of their visual field in a certain location and nothing else. Let us say that A_1 is the phenomenal state of A, B_1 is the phenomenal state of B, and A_2 is the phenomenal state A would have had if the red flash had been in the corresponding location on the other side. What is the relationship between these states? (i) e-categoricalism: $A_1 \neq A_2$, $A_1 = B_1$ or $A_2 = B_1$ but not both; (ii) e-relationism: $A_1 = B_1$, $A_2 = B_1$, $A_1 = A_2$; (iii) Frege-Schlick view: $A_1 \neq A_2$ and there is no fact about whether $A_1 = B_1$ or $A_2 = B_1$. In principle there is also (iv) all four states are distinct, but this view is perhaps undermotivated. I am most inclined to accept (ii), but I also take (iii) seriously.

If one finds the categoricist view most intuitive here, one can motivate the alternative views by considering hypothetical two-field creatures with two separate and independent visual fields. The two fields may be quite symmetrical, so that there is no way to align one field of creature A with one field of creature B as the “same” field. Suppose both A and B have a red flash in one field, with associated phenomenal states A_1 and B_1 , and A_2 is the phenomenal state A would have had if the flash had been in the other field. Then one has the same four choices as above. One *could* embrace (i), field e-categoricism, but this seems to require alignment among fields, contra hypothesis. So it seems that one must embrace (ii)-(iv). I think that reflection on this case and on the relation to left-right cases helps to break down intuitions in favor of e-categoricism.

I conclude that there is a prima facie strong argument (from p-relationism and global supervenience) for e-relationism and against e-categoricism, and that although there are various intuitive objections to e-relationism, none are as strong as the argument for e-relationism. So I am inclined to at least tentatively accept e-relationism and reject e-categoricism.

References

- Baker, D. 2011. ‘The experience of left and right’ meets the physics of left and right. *Noûs* 46:483-498.
- Block, N. 1990. Inverted Earth. *Philosophical Perspectives* 4:53-79.
- Chalmers, D. J. 2006. Perception and the fall from Eden. In (T. Gendler and J. Hawthorne, eds.) *Perceptual Experience*. Oxford University Press.
- Chalmers, D.J. 2012. *Constructing the World*. Oxford University Press. Extended edition at <http://consc.net/books/ctw/extended.html>.
- Hofer, C. 2000. Kant’s hands and Earman’s pions: Chirality arguments for substantival space. *International Studies in the Philosophy of Science* 14:237-56.
- Hurley, S. 1998. *Consciousness in Action*. Oxford University Press.
- Kant, I. 1768. On the first ground of the distinction of regions in space. Translated in (J. Handyside, trans.) *Kant’s Inaugural Dissertation and Early Writings on Space*, pp. 19-29. Open Court, 1929.
- Lee, G. 2006. The experience of left and right. In (T. Gendler and J. Hawthorne, eds.) *Perceptual Experience*. Oxford University Press.
- Pooley, O. 2003. Handedness, parity violation, and the reality of space. In (K. Brading & E. Castellani, eds) *Symmetries in Physics: Philosophical Reflections*. Cambridge University Press.

Putnam, H. 1975. The meaning of “meaning”. In *Mind, Language, and Reality*. Cambridge University Press.

Schoettle, T. 2009. How I learned to stop worrying and love the inverted spectrum. *Pacific Philosophical Quarterly* 90:98-115.

Thompson, B.J. 2010. The spatial content of experience. *Philosophy and Phenomenological Research* 81:146-184