Brit. J. Phil. Sci. **67** (2016), 837–878

Symmetry as an Epistemic Notion (Twice Over)

Shamik Dasgupta

ABSTRACT

Symmetries in physics are a guide to reality. That much is well known. But what is less well known is why symmetry is a guide to reality. What justifies inferences that draw conclusions about reality from premises about symmetries? I argue that answering this question reveals that symmetry is an epistemic notion twice over. First, these inferences must proceed via epistemic lemmas: premises about symmetries in the first instance justify epistemic lemmas about our powers of detection, and only from those epistemic lemmas can we draw conclusions about reality. Second, in order to justify those epistemic lemmas, the notion of symmetry must be defined partly in epistemic terms.

doi:10.1093/bjps/axu049

Advance Access published on March 12, 2015

1 Symmetry-to-Reality Reasoning

1.1 A rough introduction to symmetry

Symmetries in physics are often used as a guide to reality. Let me illustrate what I mean with a simple example. Suppose that we have compelling evidence to think that the world is ultimately made up of point particles with mass. And suppose that the way they move looks (given our evidence) like they are governed by just two laws: $f = ma$ and the inverse-square gravitational force law. Call these the laws of Newtonian Gravitation (NG).

The laws of NG (like all laws) come associated with symmetries. What does this mean? We are familiar with the idea that material objects can have symmetries, for example my plate is (roughly) symmetric through straight lines that bisect its centre. But what does it mean for a law to exhibit symmetries? This is one of the central questions of this article and I will not offer an answer until near the end. But we need a rough characterization to get us started. For now, think of a symmetry of a law as a transformation on physical systems that (at a minimum) preserves the truth of the law.[1] The symmetries of NG include rigid spatial translations and rotations, rigid temporal translations, and uniform velocity boosts.

To illustrate, consider a possible closed physical system in which some particles are moving around in a way that conforms to the laws of NG. And now consider a second possible closed system just like the first, except that it is located three feet to the right of the first. In this second system the particles are moving around just as they are in the first, with one exception: at any given time, any given particle is three feet to the right of where it is in the first system. It is easy to prove that this second system also conforms to the laws of NG. After all, the masses of all its particles are exactly the same, as are their

[1] Earman ([1989]) calls this the notion of a ‘dynamical’ symmetry, but I drop the qualifier here for brevity.

distances from one another and their accelerations (at any given time). And since the laws of NG govern just these quantities, the second system must therefore conform to NG if the first one does. The operation that takes the first system as input and yields the second as output is called a rigid spatial translation, and what we have shown is that this translation preserves the truth of NG. And clearly all such translations (for arbitrary distances and directions) preserve the truth of NG too, and they are all said to be symmetries of NG.

Similarly, consider a third possible closed system just like the first, except that it is moving at 5 mph towards the north relative to the first. In this third system the particles are moving around just as they are in the first system with the one exception that at any given time any given particle is moving 5 mph faster towards the north than it is in the first system. It is easy to prove (in the same way as above) that this third system must also conform to the laws of NG if the first one did. And the operation that takes the first as input and yields the third is called a uniform velocity boost, and all such boosts (for arbitrary directions and speeds) preserve the truth of NG and are said to be symmetries of NG too.

This is (as I said) very rough as it stands. At some point we must say more about what a physical system is (a possible world, or a part of a possible world, or a model, or a set of sentences, etc). And we will see that to count as a symmetry the transformation must preserve more than just the truth of the laws. But this loose characterization communicates the rough idea well enough for now.[2]

Now, a symmetry preserves some quantities but not others. Any two possible closed systems related by a rigid translation or uniform velocity boost clearly agree on all facts about the distances between particles at times, their relative velocities, and so on. These quantities are therefore called invariant in NG: their values are preserved by all of NG's symmetries. But other quantities are altered by a symmetry of NG. For example, suppose that these point particles are moving around in Newtonian space so that each particle at any given time has an absolute position in that space and an absolute velocity through it. Then a (non-trivial) translation will change the absolute position of each particle, and a (non-trivial) boost will change each particle's absolute velocity. These quantities are therefore called variant in NG: if particles have these quantities then their values are altered by a symmetry of NG.

[2] I am thinking of symmetries actively (that is, as transformations on actual or possible physical systems), rather than passively (that is, as transformations on coordinates used to describe a given physical system). See (Brading and Castellani [2007], pp. 1342–3) for more on this distinction.

1.2 The symmetry-to-reality inference

The way in which symmetry is used as a guide to reality can now be easily stated. The idea is that if a putative feature is variant in laws that we have reason to think are true and complete, then this is some reason to think that the feature is not real. So, if we had reason to think that the laws of NG were true and complete then the idea is that, because absolute velocity is variant in NG (that is, there are symmetries of NG that alter its values), we would also have reason to think that absolute velocity is not real, and to endorse a view of space and motion according to which there is no such thing. The resulting view might be that the point particles live in a Galilean space-time structure in which there is no notion of absolute velocity. Or it might be a more radical relationalist picture on which there is no such thing as space or space-time through which the particles are moving. But for now we need not settle on the specific metaphysics: the point is that some such view is motivated.

The line of reasoning generalizes to other physical theories too. Absolute simultaneity is (famously) variant in the laws of the special theory of relativity. And it is commonplace to think that, insofar as we consider those laws to be true and complete, we should think that there is no such thing as absolute simultaneity (perhaps by endorsing the view that space-time instantiates a Minkowski structure).

From cases like these we can abstract the form of these 'symmetry-to-reality' inferences as follows:

(1) Laws L are the complete laws of motion governing our world.

(2) Feature X is variant in L.

(C) Therefore, X is not real.

This form of inference is ubiquitous. Earman's principle SP1 ([1989], p. 46)—according to which it is a condition of adequacy of one's metaphysics of space-time that it not contain variant features—is (in effect) an endorsement of this inference, and he puts his principle to work throughout his book. North ([2009], p. 64) calls this form of inference a 'methodological principle' in physics, and she uses it to investigate the metaphysical implications of Hamiltonian and Lagrangian formulations of classical mechanics. And Baker ([2010], p. 1157) writes that this use of symmetries 'has many metaphysical applications' that he then explores.

The inference is also implicit in the motivation behind recent structuralist theses about space-time. For example, consider recent discussions of diffeomorphism invariance in general relativity. These are symmetries that alter facts about which particular regions of the manifold instantiate which parts of the fields. And many have inferred from these symmetries that there must be a sense in which the particular regions—when considered as independent of the fields—are not real. On one reconstruction, this is just the symmetry-to-reality inference at work.[3]

So far I have focused on symmetries like boosts and translations. But it is also natural to consider other transformations—for example, uniform multiplications of mass (perhaps with an associated alteration of the gravitational constant or spatial distances). If such transformations can be shown to be symmetries of NG in the same sense that boosts and translations are, the symmetry-to-reality inference would have us endorse a view of mass on which only the mass relationships that are invariant are real and any further notion of intrinsic mass is unreal.[4] Or consider parity inversions, namely, transformations that flip left and right. Again, if these can be shown to be symmetries of NG in the same sense that boosts and translations are, the symmetry-to-reality inference would have us conclude that there is no real difference between right- and left-handed gloves, other than their being incongruent with one another.[5]

The potential scope of symmetry-to-reality reasoning is therefore easy to appreciate.

1.3 Two questions

Here I ask two (related) questions about this reasoning. First, how is it to be justified? I ask this in the spirit of reconstructive epistemology, much as one might ask why we are justified in trusting the testimony of others, or why induction and *modus ponens* are justified. One could ask these questions as a sceptic, but one need not. For even if we assume that they are good inferences, the question remains as to why they are good.[6] The question is not about the sociological history of these inferences or the psychological processes involved when making them. It is rather an epistemic question of justification. Why do premises about symmetries justify metaphysical conclusions about what is real? What is wrong, epistemically speaking, with favouring a Newtonian

[3] I offer one reconstruction along these lines in (Dasgupta [2011]). Hoefer ([1996]) and Ladyman and Ross ([2007]) also offer reconstructions along similar lines. Admittedly, sometimes the inference here goes via issues of determinism, as in the presentation of the hole argument in (Earman and Norton [1987]). But according to the reconstruction I have in mind, the detour via determinism is inessential to the matter.

[4] I develop this line of argument in (Dasgupta [2013]). A related example concerns gauge invariance in classical electromagnetism.

[5] See (Pooley [2003]) for a discussion of this issue.

[6] Belot ([2013]) has recently argued that they are not justified. But (as I discuss later) we disagree less than might be apparent. For what he argues is that if symmetry is defined in various ways then the inference is unjustified, and I agree with that.

view of space over a Galilean one? What epistemic mistake would one be making?

It is not immediately obvious. Someone who believes that the laws of NG are true and complete but who favours the Newtonian view of space over a Galilean one cannot be accused of logically inconsistent belief. This is just to say that these inferences are not valid. But they are often good, just as inductive inferences are invalid but (sometimes) good. I want to know what makes them good, when they are good.

One might say with Earman that it is a condition of adequacy on a space-time theory that it does not contain features that vary under the symmetries of the laws. But this just restates our data that symmetry-to-reality inferences are (sometimes) good. What I want to know is why this should be a condition of adequacy on a space-time theory.

One answer might be that there is a basic epistemic norm favouring metaphysical theses that dispense with variant features, where by 'basic' I mean that it cannot be justified in terms of further epistemic norms. Examples of basic epistemic norms might be perception (trust your senses), induction, and *modus ponens*. It is precisely because these norms are basic that the classic sceptical problems arise: if the sceptic asks why we should expect the future to resemble the past it is notoriously difficult to say anything that does not sound question-begging! But it is implausible that symmetry-to-reality inferences are epistemically basic in the same way that induction arguably is. Rather, one suspects that if a symmetry-to-reality inference is good, then that is because (i) variant features have some property (such as being redundant) and (ii) there is some more basic norm advising us to dispense with features with this property. Our question, then, is what this property is.

The question becomes pressing when one looks at what is meant by 'symmetry'. For the notion is often defined in purely formal, mathematical terms, so that whether a given transformation is a symmetry of a given set of laws depends just on the formal and mathematical features of those laws and their models. But why should those features of the laws have anything to do with metaphysics, with what's real? It is not obvious—at least not at first glance.

There is a question, then, as to what 'symmetry' could mean, such that the symmetry-to-reality inference is justified. That is my second question.

1.4 Two answers

I will argue for epistemic answers to both questions. With respect to the first question, I argue (in Sections 2–4) that we have reason to think that variant features are not real because (i) they are undetectable, and (ii) there is a more basic Occamist norm advising us to dispense with undetectable structure. As a result, the symmetry-to-reality inference proceeds via an epistemic lemma concerning detectability, as follows:

(1) Laws L are the complete laws of motion governing our world.

(2) Feature X is variant in L.

(3) Therefore, X is undetectable (from (1) and (2)).

(C) Therefore, X is not real (from (3) and an Occamist norm that we dispense with undetectable structure).

The second question, then, becomes the question of what 'symmetry' could mean, such that the inference to (3) is justified. I will argue (in Sections 5 and 6) that the standard definitions in formal and mathematical terms leave that step unjustifiable, but that it is justifiable if 'symmetry' is defined in broadly epistemic terms.

That is why I think that symmetry is an epistemic notion twice over: once with respect to its definition, and once with respect to its role in metaphysical reasoning.

1.5 Preliminary clarifications

Before I argue for this, two points of clarification might be helpful. First, I have so far ignored a number of distinctions, including the distinction between global and local symmetries, between internal and external symmetries, and between continuous and discrete symmetries; see (Brading and Castellani [2007]; Kosso [2000]) for helpful characterizations of these distinctions. This is because, *prima facie*, the symmetry-to-reality inference appears to apply generally. This is evidenced by the examples in Section 1.2: uniform boosts are global, external, and continuous symmetries; diffeomorphic shifts are local symmetries; multiplications of mass are internal symmetries; and parity inversion is a discrete symmetry. It may be that in the final analysis the inference turns out to be reasonable only for some specific class of symmetries. But that would have to be shown on the basis of an account of how the inference works, so for the most part I will continue to set these distinctions aside. (I suspect that, on the account I develop here, the inference is indeed blind to these distinctions.)[7]

Still, my discussion will mostly focus on the case of NG and its (global, external, continuous) symmetries, like boosts and translations. It is sometimes thought that the symmetry-to-reality inference is easy to justify in this simple

[7] For the same reason, I set aside the distinction between symmetries that can be physically implemented on sub-systems (such as uniform boosts) and those that cannot (such as diffeomorphic shifts). Brown and Sypel ([1995]) discuss this distinction in some detail.

case and that the more interesting question is whether it generalizes to other kinds of symmetries and other theories.[8] But I believe that even in this simple case the inference is more involved than is sometimes recognized. So I focus on the simple case, in the hope that by clarifying it we will be better equipped to see whether and how it generalizes.

Second, I will remain neutral as to how one might justify (1) and (2) in the first place—that is, how one might discover that a certain transformation is a symmetry of the actual laws. On a 'laws first' approach, one first discovers what the laws are and then works out what their symmetries are. Alternatively, on a 'symmetry first' approach, one starts with a principle to the effect that any law worthy of our credence must have certain symmetries, and one uses this as a guide to what the laws must be.[9] This principle that the laws must have certain symmetries could be *a priori*, or it might be justified on the basis of observation.[10] But there is no need to restrict our focus to either approach. Our question is: once we believe that a feature varies under the symmetries of the laws (for either of these reasons), what can we reasonably infer from that?

All I assume is that the inference does not go in a circle—that is, that (1) and (2) are not justified by a prior belief that the feature X is unreal. Thus my question (more precisely) is: once we believe that X varies under the symmetries of the laws (where this is not justified by a prior belief that X is unreal), why is it then reasonable to infer that X is unreal?

2 Against Redundancy

2.1 Redundancy

As I said, my first claim is that the symmetry-to-reality inference proceeds via an epistemic lemma concerning undetectability. To argue for this, let me start by objecting (in this section and the next) to two alternative reconstructions of the inference. On the first alternative, the inference goes via a claim about redundancy. On the second, it goes via a claim about objectivity.

The first alternative is perhaps the most common. On this view, we have reason to think that variant features are not real because (i) variant features

8 For example, Belot ([2013]) argues that the inference does not generalize, but he says relatively little about how the inference is to be justified in the simple cases.

9 An example of this is Einstein's use of the principle of relativity in the derivation of his special theory. Brading and Castellani ([2007], pp. 1346–9) give a brief history of this use of symmetry. To be clear, this question of whether laws or symmetries come first in the order of justification is not the question of which come first in the order of explanation. Lange ([2007]) discusses this latter question.

10 It is an interesting question which kinds of symmetries can be observed and which cannot, but there is no need to settle it here. For some recent work on this question see (Brading and Brown [2004]; Greaves and Wallace [2014]; Kosso [2000]; Healey [2009]).

are redundant, and (ii) there is a more basic Occamist norm advising us to dispense with redundant features. The inference then looks like this:

(1) Laws L are the complete laws of motion governing our world.

(2) Feature X is variant in L.

(3*) Therefore, X is redundant (from (1) and (2)).

(C) Therefore, X is not real (from (3*) and an Occamist norm that we dispense with redundant structure).

To be clear, 'redundant' and its synonyms (like 'superfluous') have many uses, including (perhaps) epistemic uses on which they mean something close to 'undetectable'. But read like that, this reconstruction collapses into my own. Since my aim in this section is to explore alternative reconstructions, I will focus on non-epistemic readings of the term. Thus according to this reconstruction, we can draw metaphysical conclusions from symmetries without the detour through epistemic claims about what is undetectable. I suspect that this is the standard view (if any view deserves that title): as we will see, it appears to be endorsed by a number of theorists including Baker ([2010]), Earman ([1989]), and North ([2009]), to name just three.[11]

But I do not think this reconstruction works. For what does 'redundant' mean? In order for this reconstruction to work, there must be a non-epistemic meaning of 'redundant' such that (i) being redundant in that sense follows from being variant (so that the inference to (3*) is justified), and (ii) being redundant in that same sense gives us reason to think that the feature is not real (so that the inference from (3*) to (C) is justified). I will argue that there is no meaning that plays both these roles.

It might seem impossible to argue for this until we determine what 'symmetry' means, and hence what 'variant' means, for until then we will not know what the real content of (2) is. But we can proceed by thinking just about absolute velocity. For uniform velocity boosts are paradigm examples of symmetries of NG, so whatever 'symmetry' means, absolute velocity had better turn out to be variant in NG. So if this reconstruction of the inference is to work there must be a non-epistemic meaning of 'redundant' on which (i) absolute velocity is redundant in NG, and (ii) being redundant gives us reason to think that the feature unreal. I will argue that there is no meaning that plays both these roles.

[11] As an anonymous referee pointed out, this exegetical claim needs argument. Part of the problem is that if 'redundant' has an epistemic interpretation then a theorist who appears to use this reconstruction may in fact have my epistemic reconstruction in mind. A full discussion of this issue would take us too far from the main thread, but the quotations later in this section suggest that these authors had non-epistemic interpretations in mind.

2.2 Is absolute velocity redundant?

Let us then canvass some different readings of 'redundant'. On one reading, a feature is redundant if and only if it makes no difference to how a system will evolve over time. Baker ([2010]) appears to use this reading when defending this version of the symmetry-to-reality inference. He says that we should dispense with variant features because they are not difference-makers. And something is a difference-maker, he says, if it makes 'some difference in how the state of a (physically possible) world evolves in time' (p. 1159).

Now, it may be that if a feature is redundant in this sense then we have reason to think that it is unreal. But the trouble is that variant features are not necessarily redundant in this sense: even if the laws of NG are true and complete, absolute velocity is not redundant in this sense of the term.

Why might one think that absolute velocity is redundant in this sense? Consider a physical system composed of point particles with masses moving around in Newtonian space in accordance with NG. At t_1 it is in a certain state, and the laws of NG imply that by a later time, t_2, it will have evolved into another state. Now consider all the interparticle distances at t_2. Symmetry considerations show that facts about absolute velocity in the earlier time t_1 make no difference to those later interparticle distances. After all, a uniformly boosted system is one in which (i) the earlier absolute velocities are all uniformly different, (ii) the later interparticle distances are the same, and yet (iii) the laws of NG still obtain (by the fact that uniform boosts are symmetries of NG). It follows that the absolute velocities of the particles at t_1 can vary uniformly in uncountably many ways and the laws of NG will still imply that the particles will end up standing in exactly the same interparticle distances at t_2 as they actually do. In this sense, the absolute velocities at t_1 make no difference to future facts about interparticle distances. My guess is that this line of thought is what lies behind Baker's suggestion.

This line of thought is correct, as far as it goes. But it does not go far enough, for the absolute velocities at t_1 do make a difference to facts about absolute velocity at t_2! In a boosted system, the absolute velocities at t_1 are all different and the absolute velocities at t_2 are all correspondingly different too. So if there is such a thing as absolute velocity, it is a difference-maker even though it is variant. So in this sense of 'redundant', a feature's being variant does not imply that it is redundant.[12]

Of course, one might conclude on the basis of a symmetry-to-reality inference that absolute velocity is not real, and it would then be true that it is not a difference-maker. But that would beg the current question: we are trying to

[12] A similar point is made by Sklar ([1977]).

establish how the symmetry-to-reality inference works in the first place, so we cannot assume that absolute velocity is unreal.

So much for difference-making. A second (and related) idea is that something is redundant if and only if it is dispensable from explanations of physical phenomena. Now, it may be that if something is redundant in this sense then we have reason to think that it is not real. But the point just made shows that variant features are not necessarily redundant in this sense: the absolute velocities of the particles at t_1 will be vital to explaining their absolute velocities at t_2 precisely because if their current velocities were different then their later velocities would differ too.

Of course, if absolute velocity were undetectable then it would be dispensable from all explanations of detectable phenomena. And perhaps that would be reason to think that it is unreal. Thus one might suggest that something is redundant if and only if it is dispensable from all explanations of detectable phenomena. But to show that absolute velocity is redundant in this sense we would first have to show that it is undetectable, and we are currently exploring justifications of the symmetry-to-reality inference that do not proceed via this epistemic lemma. (Put differently, if we define 'redundant' this way then the current reconstruction in terms of redundancy is not an alternative to my own epistemic reconstruction.)

A third suggestion is that a feature is redundant if and only if its values are arbitrary.[13] The idea is that because absolute velocity varies under the symmetries of NG, there is no explanation as to why the absolute velocities of things are as they actually are, rather than some boosted variation.

But the suggestion is confused. If the idea is that the absolute velocity of a given particle can never be explained, then the idea is false: we just saw that its absolute velocity at a given time will be explained (in NG) by its absolute velocity at prior times (and the forces operating on it). Perhaps the thought is that the absolute velocities at the initial condition are unexplained and arbitrary. If so, the charge is correct but would apply to the values of all features, whether variant or invariant, and so cannot be that property of variant features that would justify dispensing with them.

2.3 Some redundancies

We have discussed senses of 'redundant' on which absolute velocity is not redundant. There are other senses in which it is redundant, but unfortunately they are not senses that give us reason to dispense with it.

One idea is that a feature is redundant to some laws if and only if it is not needed to support those laws, where this means that one can formulate the

[13] This sort of idea obviously has its roots in the principle of sufficient reason.

laws without assuming that the feature is real. Perhaps this is what Earman had in mind when he wrote:

> The motivation [for the symmetry-to-reality inference] derives from combining a particular conception of the main function of laws of motion with an argument that makes use of Occam's razor. Laws of motion, at least as they relate to particles, serve to pick out a class of allowable or dynamically possible trajectories. If [there are variant features], the same set of trajectories can be picked out by the laws working in the setting of a weaker space-time structure. The theory that [posits variant features] is thus *using more space-time structure than is needed to support the laws*. (Earman [1989], p. 46, my emphasis)

It may also have been what North had in mind when she justified the symmetry-to-reality inference by saying that 'we should go with the *least structure*. We should not posit structure beyond that which is indicated by the fundamental dynamical laws' ([2009], p. 64). And it may have been what Belot had in mind when he wrote, 'Why was it (in hindsight, in one sense) a mistake for Newton to postulate absolute space? Because he thereby postulated *more spacetime structure than was required for his dynamics*' ([2013], p. 328).

It is uncontroversial that absolute velocity is redundant to NG in this sense: it is possible to formulate the laws as governing the trajectories of particles through a Galilean space-time structure in which there is no such thing as absolute velocity.[14] It is not immediately obvious that in all cases in which a given feature is variant in some laws, it is redundant to those laws in this sense. But let us grant that this is so for the sake of argument.

The trouble is that it is not clear why a feature's being redundant in this sense is reason to think that it is not real. Why should the fact that I can describe the laws of motion without mentioning a given feature constitute evidence that the feature is unreal?

This point is most perspicuous if one considers a physical theory that can be divided into two components. The first component describes the fundamental physical ontology of the world, the kind of stuff that the physical world is made up of. It might describe the kinds of particles or fields there are, the kind of space-time structure they live in, and so on. Call this component the theory's 'metaphysics'. And the second component describes how that aforementioned physical stuff behaves. These are the theory's 'laws'. The remarkable thing about the symmetry-to-reality inference is that it

[14] I am skipping over a difficulty here. For one might argue that the laws when formulated over a Galilean space-time structure are different than when formulated over a Newtonian space. After all, the quantity of acceleration will be defined differently in each case. So if the statement in the text is true, that is because we are using a notion of sameness of law that would count them as being the same law even in light of these differences. But it would be distracting to elaborate on this here.

uses facts about one half of your theory (the laws) to inform the other half (the metaphysics). But it is highly unobvious why this is reasonable—after all, they are different parts of the theory! Earman's idea is that if you can write the laws without mentioning some feature in the metaphysics, then that is reason to think that the feature is unreal. But why should this be so? Why think that every real feature must be mentioned when writing down one half of your theory?

Keep in mind that even though one can formulate NG without mentioning absolute velocity, absolute velocity is nonetheless a difference-maker and is indispensable to explanations of physical phenomena in an NG system, in the senses discussed above. Why then is the mere fact that one can formulate the laws without mentioning it a reason to think that it is unreal?

The same point applies to another sense of 'redundancy'. For one might say that a feature is redundant to some laws if and only if its values are irrelevant to whether the laws obtain. Absolute velocity is redundant to NG in this sense: the particular absolute velocities of things are irrelevant to whether the laws of NG obtain, in the sense that we can uniformly change them however we wish and the laws will remain true. Moreover this follows from the fact that absolute velocity is variant in NG. Indeed, it may be that this is the sense of 'redundancy' that Earman had in mind in the above quote.

But the trouble again is that it is not at all clear why being redundant in this sense is reason to think that a feature is unreal. Remember, a feature can be redundant in this sense and yet (like absolute velocity) be a difference-maker and be indispensable to explanations of physical phenomena. All redundancy means, on this interpretation, is that the particular values of a feature from one half of our theory (the metaphysics) are irrelevant to whether the other half of the theory (the laws) is true. But there is no obvious reason to think that every categorical aspect of the world must be relevant, in this sense, to whether the laws obtain.

For these reasons, I do not think that this non-epistemic reconstruction of the symmetry-to-detection inference, in terms of redundancy, succeeds. Perhaps its initial attraction rests on an equivocation between different notions of redundancy. It is true that absolute velocity is redundant in NG in the two senses just discussed: it is not needed to formulate the laws and its values are irrelevant to whether NG is true. And it may be true that being redundant in the earlier senses—not being a difference-maker, being dispensable from explanations of physical phenomena—is reason to think that a feature is unreal. But the senses of redundancy are not the same.

Now, there is a sense of 'redundant' on which (i) absolute velocity is redundant in NG, and (ii) its being redundant implies that it is not real. It is the sense in which a feature is redundant if and only if models that differ only in a uniform transformation of its values must represent the same possible world.

Now, it may be that models related by a uniform velocity boost represent the same possible world, and so it may be that absolute velocity is redundant in this sense. But the problem with this sense of 'redundant' for our purposes should be obvious: we would only be justified in believing that uniform velocity boosts must represent the same possible world if we already had reason to think that absolute velocity is not real. So the claim that velocity is redundant in this sense can hardly be used as a non-question-begging premise in an argument that it is not real.

These points are all straightforward, so why are they often ignored? One reason might be the literature's focus on technical formulations of the matter, which tends to obscure the important epistemic questions. It is easy to show that physical situations differing in a uniform velocity boost are related formally in various ways, and that the variant features are therefore redundant in some formal sense. But this is not enough. The question remains as to why a feature's being redundant in that formal sense is reason to believe that it is unreal. As we have seen, it is surprisingly difficult to answer the question.

Indeed there is an independent consideration that raises a doubt about any reconstruction of the symmetry-to-reality inference in terms of redundancy. If we could see or detect absolute velocity, we should believe that it is real regardless of whether it is redundant in any of the above senses. So a necessary condition on our reasonably believing that it is unreal is that we do not believe that it is detectable. And from there one might argue that a necessary condition on our reasonably believing that it is unreal is that we believe that it is undetectable. If that is right then it is hard to see how any non-epistemic reconstruction of the inference in terms of redundancy could work.

3 Against Objectivity

Before discussing my own reconstruction of the symmetry-to-reality inference, let me set aside another alternative: one that proceeds via a claim about objectivity.[15]

The idea that there is a link between symmetry and objectivity has a distinguished history. Weyl ([1952], p. 132) famously made this link when he said that 'objectivity means invariance'. And it was recently taken up by Nozick ([2001], p. 76), who proposed that 'an objective fact is invariant under various transformations'.[16] The idea is that something is objective if and only if it is invariant—an equation of invariance and objectivity.

Suppose one accepts this equation. Then one might infer from something's being variant to its being non-objective. And one might then take this to

[15] Thanks to an anonymous referee for suggesting this alternative to me.

[16] For further discussion of this idea, see (Debs and Redhead [2007]; Kosso [2003]).

suggest that it is not real. Thus one might reconstruct the symmetry-to-reality inference as follows:[17]

(1) Laws L are the complete laws of motion governing our world.

(2) Feature X is variant in L.

(3**) Therefore, X is not objective (from (1) and (2)).

(C) Therefore, X is not real (from (3**) and an Occamist norm that we dispense with structure that is not objective).

Unfortunately, this reconstruction fails for two reasons: the inferences to (3**) and to (C) are both objectionable.

This is clear once we clarify the relevant notion of objectivity. Roughly speaking, the notion in play is that of 'perspectival independence', of being the same from all perspectives so that there is agreement amongst all observers.[18] To see why an equation of invariance and objectivity so understood might make sense, it helps to think of symmetries 'passively'. Until now we have been thinking of symmetries 'actively', as ways of altering physical systems. In contrast, when we think of a symmetry passively 'we re-describe *the same physical evolution* in two different coordinate systems', as Brading and Castellani ([2007], p. 1342) put it. We take a description of a physical process in terms of a given coordinate system, and we think of (say) a boost as a re-description of the same physical process in terms of a boosted coordinate system. The two coordinate systems will disagree on the velocities of bodies, and so velocity is not invariant between these two descriptions. And so velocity is not an objective quantity: it depends on a choice of coordinate systems, and so is not independent of perspective. To take another example (used by Nozick), consider the Centigrade and Fahrenheit temperature scales. Famously, ratios are not invariant between these two scales: $100°C = 212°F$ and $50°C = 122°F$, yet 212 is not twice 122. So the ratio between two temperatures in a given scale is not objective: it depends on an (arbitrary) choice of scales.

Having clarified the notion of objectivity in play, we can now see that the inference from (3**) to (C) is unwarranted. To see this, notice that (just like velocity) acceleration is not objective in the relevant sense of the term: two coordinate systems that are accelerating relative to one another will disagree about the rate of acceleration of a given body—in particular, each coordinate

[17] To be clear, I do not know whether Weyl or Nozick would endorse this, and I have not seen it suggested in the literature. But it is an intriguing idea.

[18] This is (more or less) what Nozick means by 'objectivity'. Is this what Weyl meant? It is not so clear, and since I am no scholar I will not try to decide the matter here. Still, Daston and Galison ([2007], Chapter 5) argue that this notion of objectivity played an important role in scientific theorizing in the early-twentieth century.

system is at rest according to itself but accelerating according to the other. So, just like velocity, acceleration is not independent of perspective and so is not objective in the relevant sense.[19] But if NG is true then acceleration is arguably real (for bucket-argument-type considerations). So the inference from (3**) to (C) is bad.

It might sound odd to call acceleration real but not objective, but only if one forgets that the two notions of reality and objectivity are very different. The former is a metaphysical notion concerning what the world is made up of, while the latter is a more epistemic notion of agreement between different perspectives (scales, coordinate systems, and so on).[20] So it is no surprise that the two notions come apart. One could try to force the two notions into coextension by requiring that the laws be objective, for example that the equations are invariant under all coordinate transformations.[21] But the symmetry-to-reality inference is a good inference even when we consider laws that are not invariant in this way, such as classical NG. What justifies the inference in these cases? Not, as we just saw, an appeal to objectivity.

Moreover, we can now see that if the inference to (3**) is justified by the equation of invariance and objectivity, then it is based on an equivocation. For what are the relevant symmetries, invariance under which is equated with objectivity? Not symmetries of the laws. For acceleration is invariant under the symmetries of NG; but, as we have just seen, acceleration is not objective in the relevant sense (even if NG is true). Those who equate invariance with objectivity must therefore specify the relevant symmetries differently. Indeed, Nozick ([2001], Chapter 2) explicitly takes up this task. So, when he equates objectivity with invariance, his notion of invariance is different than the one I use in premise (2).

To be clear, none of this is an objection to Weyl or Nozick's equation of objectivity and invariance (in some sense of the term). The point is just that this equation does not help justify the symmetry-to-reality inference.

[19] It is true that $f = ma$ will not be true when evaluated relative to at least one of those coordinate systems. But this does not contradict the point that acceleration is not objective in the above sense.

[20] Daston and Galison ([2007], Chapter 5) discuss the difference between these two notions and the role of each in the early-twentieth century.

[21] This would be to use a symmetry principle as a guide to what the laws must be; see Section 1.5. This idea that the equations should be invariant under all coordinate transformations is related to the requirement of general covariance. However, the precise meaning of general covariance, and its role as a symmetry principle, is notoriously controversial. See (Howard [1999]) and (Ryckman [2005], especially Chapter 2), for some discussion of these issues.

4 From Symmetry to Detection

4.1 The epistemic approach

So far I have discussed two alternative reconstructions of the symmetry-to-reality inference and found them both wanting. Instead, I believe that the inference must proceed via an epistemic claim of undetectability, as outlined earlier:

(1) Laws L are the complete laws of motion governing our world.

(2) Feature X is variant in L.

(3) Therefore, X is undetectable (from (1) and (2)).

(C) Therefore, X is not real (from (3) and an Occamist norm that we dispense with undetectable structure).

But now I must say why this is a good inference.

4.2 The Occamist norm

I will focus on the inference from (1) and (2) to (3), but let me say something briefly about the inference from (3) to (C). Verificationists might say that it is analytic, that it is part of the meaning of 'reality' that everything real is detectable. But we need not take such a strong line. All we need (roughly speaking) is a plausible Occamist norm to the effect that positing undetectable structure is an epistemic vice, in the sense that, all else being equal (or near enough), we should prefer theories that do not posit undetectable structure. (This will be qualified as we go along, but this will do for now.)

The plausibility of this norm depends on what is meant by 'undetectable'. If the term was used to include anything that we cannot see with the naked eye, the norm would recommend that we become radical scientific anti-realists and dispense with so-called theoretical entities such as electrons. But that is not how I use the term. Instead, something is undetectable in my sense if, roughly speaking, it is physically impossible for it to have an impact on our senses. Electrons are detectable in this sense because there are physically possible processes, such as those that occur in particle accelerators, by which the presence of an electron can be made to have an impact on our senses via its impact on (say) the movement of a dial or an image produced on a computer screen. In contrast, the idea is that if the laws of NG are true and complete then it is physically impossible for the absolute velocity of any given body to have an effect on our senses.

So understood, the motivation behind this Occamist norm is forthcoming. Roughly speaking, we should think that a given physical feature is real only if we have empirical evidence that it is real. If we can show that the feature is undetectable, then we will have shown that we do not (and cannot) have empirical evidence in the form of observations or measurements of it. It remains possible that dispensing with the feature yields a theory that has too many other vices to warrant belief, such as being too inelegant or complex. In that case, we would have empirical evidence of sorts that the feature is real, in the sense that our all-things-considered best empirically confirmed theory implies that it is real. But that is a situation in which all else is not equal. When that kind of situation does not arise—when the theories that dispense with the feature do not have those other epistemic vices—we will have no empirical evidence to think that the feature is real after all. And so in those situations we should endorse those theories that dispense with it. Which is what our Occamist norm states.[22]

Note, then, that if this is how symmetry-to-reality inferences work, we can draw the conclusion of the inference only when we have the alternative theory in hand and have shown that all else is equal. This explains why it was rational for Newton to believe in absolute velocity even though he knew that it was variant in NG and undetectable. The reason this was rational for him was that he had no good alternative theory to hand. He had good reason (his bucket argument) to think that relationalism was not empirically adequate. And relationalism was the only alternative view he knew of (he was not aware of Galilean space-time structures in which there is a well-defined feature of absolute acceleration, as required by his bucket argument, but no absolute velocity). So for Newton, all else was not equal and he was rational to believe in absolute velocity.[23]

4.3 From symmetry to detection

Let us now focus on the move from (1) and (2) to (3). How can we argue, just on the basis that a given feature is variant in what we take to be the true and complete laws of motion, that it is undetectable? Let us start simple, again with

[22] I am being deliberately non-committal about how to understand empirical evidence. The argument in the last paragraph can be cashed out in a Bayesian framework, or in an abductive framework.

[23] There is a lesson here for contemporary structuralists, such as Ladyman and Ross ([2007], p. 130), who take the fact that diffeomorphisms are symmetries of general relativity to suggest that '[t]here are no things. Structure is all there is'. For it is not enough to note that individual points of the manifold are variant features and declare that they are therefore not real. That would be analogous to Newton declaring that there is no such thing as absolute velocity without a genuine alternative theory in hand, a move that we would rightly have regarded with suspicion. To motivate structuralism, one must present a clear theory of the fundamental structure of the material world without making reference to regions of the manifold, a theory that does well on other theoretical virtues such as simplicity, elegance, and so on. But contemporary structuralists tend not to present such a theory. I say more about this, and make a start at presenting such a theory, in (Dasgupta [2011]).

the case of absolute velocity. How can we establish that it is undetectable, merely on the assumption that it is variant?

One might think that we can establish this without the assumption. Any rail traveller can testify that events inside a carriage look the same regardless of whether the train is at rest or in uniform (unaccelerated) motion, and one might infer from this that absolute velocity is undetectable. But this is inconclusive. Even if ordinary events on trains do not distinguish different inertial states of motion, the possibility remains that (with enough funding) we could build a fancy device that is sensitive to different inertial states of motion and use it to measure absolute velocity. For absolute velocity to be undetectable (in my sense), this must be impossible.

So how can we show that this is impossible, merely on the assumption that absolute velocity is variant? It would be easy to argue from authority. Feynman himself says that 'the laws of Newton are of the same form in a moving system as in a stationary system, and therefore *it is impossible* to tell, by making mechanical experiments, whether the system is moving or not' ([1963], p. 15, my emphasis). And Earman says that because 'Newton's laws of motion and gravitation have (Gal) as their dynamical symmetries, *no feature* of the lawlike behavior of gravitating bodies can be used to distinguish an absolute frame: in this sense, absolute space is unobservable' ([1989], p. 48, my emphasis). It is clear from the respective contexts that both authors are inferring from a premise about the symmetries of NG to the conclusion that absolute velocity is undetectable in my sense. And yet neither Feynman nor Earman says anything to justify the inference. So is a justification available?

As it turns out, there is an argument that does the trick.[24] At a first approximation, it goes like this: Presumably, a necessary condition on absolute velocity being detectable is that there is some physically possible process that, when initiated to measure the absolute velocity of a given body at t_0, will generate a reading at t_1—an image on a computer screen, say, or the position of a needle—that indicates what that body's velocity was at t_0. Moreover, the outcome that would be produced if the body were travelling at one absolute velocity at t_0 must be discernibly different from the outcome that would be produced if it had a different absolute velocity at t_0, on pain of our not being able to tell what absolute velocity a given outcome indicates.[25] So, if we simply wanted to measure whether a given body was in a state of absolute rest or absolute motion, the process would need to produce one outcome if the body

[24] Roberts ([2008]) presented this kind of argument beautifully, and as far as I know was the first to do so in print. I have heard David Albert and Tim Maudlin both give variations of this argument in presentations.

[25] At least, that is the ideal: in practice, we do not mind if the outcomes produced by velocities differing only by some tiny amount are indiscernible. More accurately, then, what we require is that the measurement outcomes are discernible when the velocities differ by more than some amount X, in which case we say that the process measures velocity up to an accuracy of X.

was at rest at t_0—for example, an inscription of 'at rest'—and a discernibly different outcome if the body was moving t_0—an inscription of 'moving', say. Finally, since the process is a physical one, the outcome produced will depend on the physical laws governing it. Putting this all together, we can say that absolute velocity is detectable only if there is a physically possible device that at a given time, t_0, has two properties:

(i) according to the laws, it will produce one outcome at a later time, t_1, if and only if it was presented with a body at rest at t_0; and

(ii) according to the laws, it will produce a discernibly different outcome at t_1 if and only if it was presented with a body that was moving at t_0.

But if the laws of NG are true and complete, we can argue that it is physically impossible for a device to have both properties. For suppose I take a device with the first property and present it with a body at rest at t_0, and it gives some outcome. For the sake of definiteness, one can imagine that it displays 'at rest' on a screen at t_1. We can show that it does not have the second property by considering a world W just like ours, with the one exception that it has been subjected to a uniform velocity boost, say, 5 mph to the north. We know three things about W. First, it is a world in which the device is presented with a moving body at t_0 (rather than a body at rest). Second, we know (by construction) that the relative positions of all bodies at all times are the same in W as they actually are. So, if the device actually displayed 'at rest' at t_1, it displays 'at rest' at t_1 in this boosted world too. More generally, it appears that a world differing only in an absolute velocity boost, and thereby agreeing with the actual world in respect of the relative positions of all material bodies at all times, would be indiscernible from the actual world. So no matter what outcome the device gives in the actual world, we know that it gives an indiscernible outcome in the boosted world. But because boosts are symmetries of NG, the third thing we know about W is that it is a world in which the laws of NG obtain. So the behaviour of the device in this boosted world represents how it behaves according to the laws of NG; hence it does not have the second property listed above. QED.[26]

[26] As I mentioned in Section 1.5, it may be that in some cases we can observe that the actual laws have a certain symmetry. In the case of boosts, Kosso ([2000]) says that we can do this by observing two systems in inertial motion relative to one another and 'observing that the outcomes of relevant experiments are the same whether the system is transformed or not' (pp. 88–9). His idea is that with enough experimentation like this we can conclude (with an ampliative/inductive inference) that the laws must be symmetric under boosts even if we are not yet sure what the exact laws are. Suppose he is right. In these cases, is the reasoning to (3) then somehow circular? No, for (3) is not used as a premise when inferring to (1) and (2); rather, (1) and (2) are justified on the basis of the observations plus ampliative reasoning. Moreover, it was not a necessary condition on making those observations or performing the ampliative inference that one was already justified in believing (3).

That seems straightforward enough in the case of velocity. But does the argument generalize to arbitrary laws and arbitrary features?

Well, the argument went through providing that three things were true of the boosted situation: First, it is a situation in which the variant feature (absolute velocity) differs. This is guaranteed by what it is to be variant. Second, it is a situation in which the laws obtain. That much is guaranteed by the fact that symmetries preserve the truth of the laws. But the third thing that needed to be true of W is that it is (in some intuitive sense to be made clear) observationally equivalent to the original world, that at least to the naked eye it looks and feels and smells exactly like the original world. This is what we were trying to show when we pointed out that W agrees on all facts about the relative positions of things and (therefore) on what is displayed on a computer screen. So what we need, if the argument is to generalize to arbitrary laws and arbitrary features, is this:

> Given any set of laws, any two systems related by a symmetry of those laws are observationally equivalent.

Is this true? It is not clear. Indeed we do not know enough about what 'symmetry' means to even start answering the question. I said at the beginning that a transformation is a symmetry of L only if it preserves the truth of L, but I also said that this is not a sufficient condition. So we are not in a position to begin asking whether the claim above is true.

The obvious strategy now is to reverse engineer. We need to ask what 'symmetry' could possibly mean such that the indent claim above is true.

5 The Meaning of 'Symmetry'

5.1 A framework

There is little agreement about the meaning of 'symmetry'. Perhaps there is nothing wrong with this: there may be many related notions each playing a different theoretical role. Our task is just to find a notion that makes sense of the symmetry-to-reality inference.

We need to start by agreeing on a framework in which to talk about symmetry. What kind of thing is a symmetry? I said that a symmetry is a transformation on physical systems. But what is a physical system? A possible world? Part of a possible world? A set-theoretic model?

There is a danger in using possible worlds or parts thereof. For suppose we think of symmetries as functions from possible worlds to possible worlds. And suppose it turns out that, necessarily, space-time has a Galilean structure in which there is no such thing as absolute velocity. Then there are no (non-trivial) uniform velocity boosts since there are no worlds that differ only in a uniform velocity boost! But surely even if that metaphysics of space-time turns out to be true we should still be able to talk about (non-trivial) uniform boosts being symmetries of NG.

So let us work instead with set-theoretic structures that are used to represent (possible or impossible) physical systems. Now, what a given set-theoretic structure represents is not intrinsic to the structure, but is due in part to how we choose to interpret it (models, like words, are our tools not our masters). So, more precisely, let us work with a domain of structures and an interpretation of each structure, where this latter is a specification of what each structure represents. Then even if there is no such thing as absolute velocity—even if there is necessarily no such thing—we can still talk of a set of structures, each representing a physical situation in which particles move through Newtonian space with well-defined absolute velocities. It is just that, so interpreted, the structures may represent incorrectly (they may even represent metaphysical impossibilities).

I said that a symmetry is a transformation on physical systems, so if we represent physical systems with structures we can represent a transformation with a function from structures to structures. But we must distinguish transformations that are taken to be generated by a recipe, from transformations that are not. To illustrate, consider the vector space V of velocities, consider some velocity v in V, and consider the function f_v on V that maps each vector v^* to the vector $v^* + v$. Then f_v naturally induces a function F_v on structures, where F_v is defined to map each structure s to a structure $F_v(s)$ in which the velocity of each thing is boosted by f_v. The function F_v, then, is a uniform velocity boost. But it is just a function: a set of ordered pairs of structures, with no explicit mention of how it was generated. To mark this, I will say that F_v is a 'bare transformation'. In contrast, I will say that the pair $<F_v, f_v>$ is a 'generated transformation'. A generated transformation is a function on structures (for example, F_v) together with a function on a property space (for example, f_v) that specifies how the function on structures is generated.[27]

We must now choose whether to work with bare or generated transformations. I choose the latter. Their virtues will emerge as we go along.

One virtue is that they allow a convenient definition of what it is for a feature to vary under a given transformation. Earlier I introduced this notion loosely, but given its central role in symmetry-to-reality reasoning

[27] The function f_v does the same thing to each velocity in V; that is, it adds v to it. So the resulting transformation F_v is global in the sense of (Brading and Castellani [2007]). But (to be clear) the framework is not restricted to global transformations since in other cases f may do something different to each element of the domain. Thus in GTR, f may be a diffeormorphism on the manifold, and the resulting transformation, F, will then be local in their sense; that is, it will have different effects at different points. I use 'property space' liberally so that it includes spaces like V, and manifolds, and sets of (say) determinate masses, so that the framework can make sense of both external and internal symmetries. And since we may allow f to be continuous or discrete, the framework can make sense of both continuous and discrete symmetries too.

we need to be more precise. First, distinguish the determinable feature (for example, absolute velocity) from its determinate values (for example, 1 kph to the north). For simplicity, let us think of a determinable as any set of perfectly determinate properties. Then the rough idea is that a transformation varies a given determinable if the transformation is generated by a (non-trivial) bijection on the determinable's determinates. More precisely: a determinable D varies under a generated transformation $<T, t>$ if and only if t is a (non-trivial) bijection on D and T is the function on structures induced by t. Thus absolute velocity varies under the above generated transformation, $<F_v, f_v>$, because f_v is a non-trivial bijection on the determinate velocities.

Our question, then, is: Which generated transformations count as symmetries of a given set of laws? There are (broadly speaking) three approaches to defining symmetry: formal approaches define it in purely formal, set-theoretic terms; ontic approaches define it with reference to certain privileged physical features; and epistemic approaches define it in epistemic terms. Most definitions in the literature are either formal or ontic, but I will argue that only an epistemic definition can make sense of the symmetry-to-reality inference.

5.2 Formal definitions

Say that a generated transformation $<T, t>$ preserves a law L if and only if for any structure, s, if s is a model of L then so is $T(s)$. Then we know that a necessary condition on a transformation being a symmetry of L is that it preserves L. Suppose that we now say that this is also a sufficient condition. This would then be a purely formal definition.

But this definition is inadequate, for it does not guarantee what is needed to underwrite the argument (from Section 4) that variant features are undetectable. For that argument to go through, it must be that given any set of laws, any two systems related by a symmetry of those laws are observationally equivalent. But there are many transformations that preserve a given set of laws but do not map systems to observational equivalents.

To see this, pick two models of NG, m and m^*, that differ in any way you like—for example, let m contain a single particle at rest and let m^* contain 10 particles in motion. And consider the function F on structures that maps m to m^* and vice versa, and which is identity on all other structures. This can be generated by considering the determinate property M that the entire system m instantiates (that is, the property of containing a single particle at rest) and the corresponding determinate property M* that the entire system m^* instantiates. The set {M, M*} is then a determinable property, so if we let f be the only non-trivial bijection on it then we have the generated transformation $\mathrm{F} = <F, f>$. F counts as a symmetry of NG according to the above definition since (by construction) it preserves NG. But it clearly does not map systems to observationally equivalent systems.

Put differently, the determinable property of containing a single particle at rest or containing ten particles in gravitational motion varies under this transformation F. So if F counts as a symmetry then, according to the reasoning outlined in Section 4, we should be able to infer that whether a system contains a single particle at rest or ten particles in gravitational motion is undetectable. Clearly, that would be incorrect. The problem is that the reasoning in Section 3 assumed that symmetries map systems to observationally equivalent systems, which F does not.[28]

If one worries that it was artificial to consider a determinable with just two determinates, one can instead consider the set WP of all world properties; that is, consider the set that contains M and M*, and indeed all the other determinate properties that an entire physical system could instantiate. Each property in WP is a fully determinate way for an entire physical system to be. Now let WP_{NG} be the subset of WP of world properties that are had by a system in which NG is true. Then there are obviously a host of non-trivial bijections, g, from WP to itself that map M to M* and vice versa, and map elements of WP_{NG} to elements of WP_{NG}. The generated function $<G, g>$ will then count as a symmetry of NG on the current definition, but will map m to m^* and vice versa, and so will not always map structures to observationally equivalent structures.

One might try to rule out these examples by insisting that a symmetry be generated from a function on a set of fundamental properties, the thought being that the world properties in WP are not fundamental. But the trouble is that they may turn out to be fundamental—indeed, that is the view of the monist who thinks that (fundamentally speaking) there is just one object, the cosmos, that has one of the properties in WP![29]

One might instead try to avoid the problem with a different kind of formal definition. The idea is to think of symmetries as transformations on (representations of) instantaneous physical states, rather than four-dimensional physical systems. Instead of requiring that the transformation preserve the laws, we would require that it commutes with the laws. This means that given

[28] Belot ([2013]) gave a similar argument against this definition of symmetry. His argument was more straightforward because he was discussing the use of symmetries that simply identifies structures related by a symmetry; since m and m^* should clearly not be identified it follows that F cannot count as a symmetry. But the symmetry-to-reality reasoning I am discussing is more complicated than his, proceeding (as it does) via lemmas about variant features being undetectable. So the reason why F is a counterexample to the current definition of 'symmetry' is a little more involved.

[29] See (Schaffer [2010]) for a contemporary defense of monism. To be clear, the point here does not depend on the truth of monism. The point is rather that the monist may legitimately wish to use the symmetry-to-reality inference to motivate various metaphysical theses, so an adequate definition of 'symmetry' should be acceptable to the monist.

any initial state X, the state produced by first evolving X in accordance with the laws and then transforming it is always the same as the state produced by first transforming X and then evolving it in accordance with the laws. Let us also require that it be a bijection on states. The resulting definition of 'symmetry' has a distinguished history dating back at least to Wigner and is used by contemporary authors such as Baker.[30] And it seems to avoid the kind objection raised above: a function like *F*, which switches two states, will not in general commute with the laws and so will not in general count as a symmetry.

Well, perhaps not in general, but we can arguably find problematic cases. Consider an NG system A in which a single particle is at rest, and a second NG system B in which two mass-less particles are at rest some distance from one another. These are both static systems, so let S_A be the state that system A is in at all times, and S_B be the state that system B is in at all times. Now consider the function *G* that maps S_A to S_B and vice versa, and is identity on all other states. This is a bijection that commutes with NG, but it is not a symmetry of NG. And the problem is not solved by requiring that a symmetry is generated from bijections on determinate properties. For *G* can be induced from a bijection on the determinate properties of entire states, in much the same way that *F* was induced by a bijection on the determinate properties of entire worlds.[31]

A formal definition of 'symmetry' must fix these problems using just the formal materials of set theory, model theory, and so on. I have not argued that this is impossible (in particular, the Wignerian approach can be developed in ways that I have not discussed here). But remember, our notion of symmetry must imply that given any set of laws, any two systems related by a symmetry of those laws will be observationally equivalent. And it is (to put it mildly) extremely hard to see how any purely formal definition could have this consequence. Absent some reason to think that 'symmetry' must be defined formally, then, let us put this approach aside.

5.3 Ontic definitions

How should we proceed? Think back to when we argued (briefly) that two situations related by a uniform velocity boost are observationally equivalent.

[30] See (Wigner [1967]; Baker [2010]).

[31] If one does not recognize the possibility of mass-less particles, one could instead consider a physical theory in which particles have charges governed by Coulomb's law. One can then let the particles in system B be massive, but let them also have charges that repel each other to exactly the same extent that their masses attract one another, so that they are at rest. Then, once again, *g* is a bijection that commutes with the laws of the theory, but is not a symmetry of the theory.

One might of course patch up the proposed definition by requiring that a symmetry be a smooth transformation on states. But this is to introduce the notion of smoothness into the definition, which is not purely formal. The resulting definition is therefore ontic, and I will discuss ontic definitions shortly.

We noted that the situations would (by construction) agree on all the relative positions of things. Perhaps it was because of this that it seemed right to say that they are observationally equivalent.

There is something compelling about this idea. As Bell ([1993], p. 166) remarked, 'in physics the only observations we must consider are position observations, if only the positions of instrument pointers'. One might object that we should also consider colour observations, like the colour of litmus paper. But still, one might think that something's colour supervenes on the relative positions of things (its parts, the light reflecting off it, and so on). So the compelling idea is not that we can observe only relative positions, but that our observations supervene on relative positions. Preserve the relative positions, Bell's idea is, and you have an observationally equivalent situation. The natural idea, then, is to require symmetries to preserve relative positions.

More precisely, say that a determinable D is preserved by a generated transformation $<T, t>$ if and only if s and $T(s)$ agree on all the determinate values of D, for any structure s. Then the natural idea is to define a symmetry of a law to be a (generated) transformation that preserves the law and also preserves relative position. It will then follow (assuming Bell's idea) that given any set of laws, any two situations related by a symmetry of those laws will be observationally equivalent, as required.

Perhaps you disagree with Bell's idea. Still, you might think that there are some other determinable features F on which our observations supervene. As a shorthand, let us say that such a collection of features F fix the data. Then the general idea is to define a symmetry of a law to be a function that preserves the law and also preserves those features F.

This general idea can be developed in two ways depending on what features F we pick. On one approach we allow F to include epistemic or observational features such as 'looking red to me' or 'appearing from my perspective to be two feet away'. This results in an epistemic definition, so put these features aside until the next section. For now, let us restrict our attention to physical features like distance, mass, charge, spin, and so on. (I will not try to define 'physical', the idea is clear enough to work with.)

The result is a definition of 'symmetry' that requires a symmetry to preserve the laws and preserve certain privileged physical features. To mark this let us call a definition of this kind an 'ontic' definition.

An ontic definition can be developed in two ways: one *de re* and the other *de dicto*. On the *de re* approach, we first pick some physical features F (that we think fix the data), and we then require a symmetry to preserve those features. On the *de dicto* approach (by contrast), we require a symmetry to preserve physical features that fix the data, whatever they happen to be. Thus the *de re* definition will say, of some features F, that a (generated) transformation $<T, t>$ is a symmetry of a law L if and only if (i) T preserves L, and (ii) T preserves those features F. And the *de dicto* definition will say that T is a symmetry of L if and only if (i) *T* preserves L, and (ii) *T* preserves all features that fix the data (whatever they are).

The *de re* definition above is just schematic. For what are the privileged physical features F that (by definition) are preserved by all symmetries? One might think there is an obvious answer: Bell's relative positions. But on second thoughts the suggestion is still schematic. For what are relative positions? Spatial distances? Or space-time intervals? Something else? A moment's thought reveals that they cannot be spatial distances, for they are not preserved by the Lorentz symmetries of special relativity.

So a proponent of the *de re* approach must propose some privileged physical features F. Earman ([1989], p. 44) suggested that a transformation <*T*, *t*> counts as a symmetry of a law L if and only if *T* preserves L and *t* is a diffeomorphism on space-time. This implies that *T* will always preserve features like smoothness and differentiability, so that (for example) if a given system, *s*, contains particles arranged in a smooth and differentiable line, *T*(*s*) also contains particles arranged in a smooth and differentiable line. So features like smoothness and differentiability are (in effect) Earman's suggestion for the privileged features F.

This *de re* ontic definition of Earman's has proved enormously popular. For example, Roberts ([2008]) uses it as his definition of 'symmetry', and Belot ([2013]) also takes it to be one of the standard definitions (he calls it the notion of a classical space-time symmetry).[32] Moreover, when Belot considers other definitions, they all turn out to require a symmetry to preserve various privileged physical features, and so they all count as ontic definitions in my sense of the term. Indeed, I suspect that the vast majority of definitions of 'symmetry' found in the literature are ontic definitions of one kind or another, though I will not pursue this exegetical claim here.

Are ontic definitions adequate? One might complain that Earman's definition is not sufficiently general: it will only count as symmetries those functions that can be generated by a diffeomorphism on space-time. This gives us (external) symmetries like boosts and shifts, but not the gauge symmetries of electrodynamics or doublings of mass (or other internal symmetries). But

[32] Why has Earman's definition been so popular? It might be due to Einstein:

> All our space-time verifications invariably amount to a determination of space-time coincidences. If, for example, events consisted merely in the motion of material points, then ultimately nothing would be observable but the meeting of two or more of these points [such as] coincidences between the hands of a clock and the points on a clock dial, the observe point-events happening at the same place at the same time. (Einstein [1916], p. 177)

Since diffeomorphisms preserve all such coincidences, the Earman definition of 'symmetry' will, if Einstein is right, map physical systems onto observationally equivalent systems.

this complaint is not weighty, since it is not hard to generalize his definition to cover these cases.

A better objection comes from Belot ([2013]), who argues that ontic definitions are not extensionally adequate: they count as symmetries transformations that vary features that we would not want to consider unreal.

But I believe that there is a more systematic problem with ontic definitions regardless of whether they are extensionally adequate. The problem is one of inferential circularity. In order to perform a symmetry-to-reality inference, I need to take what I believe to be physical laws and work out what their symmetries are. But according to an ontic definition of 'symmetry', in order to check whether a given transformation $<T, t>$ counts as a symmetry of those laws, I first need to know which physical features fix the data so that I can check whether T preserves them.[33] And the problem is that, in many cases, we discover which physical features fix the data by engaging in symmetry-to-reality reasoning! Thus on the ontic definition it is hard to see how a symmetry-to-reality inference can ever get going.

The objection rests on the claim that, in many cases, we discover which physical features fix the data by engaging in symmetry reasoning. To see why this is plausible, imagine asking someone which physical features they think would need to be preserved in order to preserve the data. 150-years ago, we would have all said that spatial distance is one such feature: if two physical systems differ (enough) with respect to the spatial distance between things, we will notice the difference. But it is crucial that this is not counted as something that fixes the data, else Lorentz transformations will not count as symmetries of the STR! And, indeed, these days we do not consider it a physical feature that fixes the data. But why not? The reason, I claim, is that we engaged in symmetry reasoning. We took it as a premise that the Lorentz transformations are symmetries of STR, and from there inferred that the features that vary under those transformations (like spatial distances) are not real, and that (therefore) what we really see when we think we are seeing spatial distance is some other quantity like space-time interval, or spatial distance relative to a frame, or what have you.

In this way our belief about whether a given feature (like spatial distance) fixes the data is based on prior beliefs as to whether the feature is real, which in turn is based on prior beliefs about what the symmetries of the physical laws are. So we cannot (on pain of inferential circularity) define symmetries to be

[33] On the *de dicto* approach this is just because I will not know which functions count as symmetries until I know which physical features fix the data, while on the *de re* approach this is because I will not know what the definition of 'symmetry' is until I know which features fix the data. But either way, it follows from any ontic definition that, in order to work out what the symmetries of a given set of laws are, I first need to know which physical features fix the data.

transformations that preserve features that fix the data (in either the *de re* or the *de dicto* sense).

To be clear, the objection is not that these ontic definitions are definitionally circular. It is just that they lead to an inferential circularity. Put differently, the objection is that they get the order of justification backwards: we often use premises about symmetries in order to work out which physical features fix the data, so we cannot at the same time define symmetries to be those operations that preserve features that fix the data.

One might reply that a *de re* ontic definition can avoid this circularity. For suppose that we just pick some physical feature and stipulate that a symmetry must (by definition) preserve it. Then we can easily know what the symmetries of a given law are without first engaging in symmetry-to-detection inferences. The above objection assumes that we require that a *de re* ontic definition appeals to features that we have reason to believe fix the data; if we relax this requirement, the objection fails. Perhaps this was what Earman intended when he offered his *de re* ontic definition requiring symmetries to preserve topological structure: perhaps there was no thought to the effect that topological structure fixes the data; perhaps he was just stipulating that the notion of symmetry must preserve it.

But if this was his idea then the problem is that the resulting definition is objectionably arbitrary. For whatever physical feature we pick as our privileged feature F, it will follow just by virtue of the resulting definition of symmetry that it is impossible to run a symmetry-to-reality inference on F and conclude that it is unreal. This is because it will be built into the very definition of 'symmetry' that F never varies under the symmetries of any law. It would clearly have been a mistake to have arbitrarily picked absolute velocity or spatial distance as such a privileged feature, for in that case uniform boosts and Lorentz transformations would by definition not count as symmetries of any law. So what is different about the differentiable structure that Earman privileges, or indeed any other supposedly privileged feature F? What is so special about them, such that they are by definition immune to being rejected as unreal on the basis of the symmetry-to-reality inference? Without an answer to this question, the definition is intolerably arbitrary.

And of course the obvious thing to say is that those privileged features singled out by a *de re* ontic defintion are privileged precisely because they fix the data. That is why they are immune to symmetry-to-teality reasoning, and why it is reasonable to build those features right into the very definition of symmetry. But, of course, if we say this then we are back with the problem of inferential circularity.

The problem of inferential circularity also arises for a final ontic definition that might be worth mentioning. One might try defining the symmetries of some laws to be those transformations that preserve the laws and that map structures to physically equivalent structures, where structures are physically equivalent if they are equally well suited to represent any given physical system.[34] But the trouble with this should be obvious, for two structures that differ (say) in a uniform velocity boost will be equally well- or ill-suited only if there is no such thing as absolute velocity. If there is such a thing as absolute velocity, then one of the structures represents a given physical system correctly if and only if the other does not. So in order to work out whether uniform velocity boosts are symmetries, on this definition, we first need to work out whether there is such a thing as absolute velocity—which (once again) gets the line of reasoning in the symmetry-to-reality inference precisely backwards.

The discussion here reveals something of an irony: the symmetry-to-reality inference is widely used, and yet influential definitions of 'symmetry' such as Earman's do not legitimate it. If there is anything going for the inference, there must be something implicit in the meaning of 'symmetry' that is ignored by these definitions.

6 Epistemic Definitions

6.1 Taking observation seriously

The implicit ingredient, I think, is that symmetries must preserve the appearances—in some sense to be made clear. If we build this into our definition of symmetry, it will follow by definition that given any set of laws, any two situations related by a symmetry of those laws are observationally equivalent. This (remember) is what we need in order to underwrite the argument (from Section 4) that variant features are undetectable. This is the epistemic approach to defining 'symmetry'.

Epistemic definitions can be developed in a number of ways. One might, for example, take as primitive a relation of observational equivalence between structures. Suppose we are in a position to know whether the relation holds between two structures without knowing anything about the underlying metaphysics of our world. Then one might simply say that a transformation $<T, t>$ is a symmetry of L if and only if (i) T preserves L, and (ii) T maps structures to observationally equivalent structures.

Still, this is not very satisfying. It is natural to think that, rather than being a primitive relation, the relation of observational equivalence holds between two structures in virtue of their intrinsic properties. So a better definition would identify the intrinsic properties that make for observational

[34] This is, more or less, Healey's notion of a 'theoretical symmetry', defined in (Healey [2009]). Though I should emphasize that he does not use this notion of symmetry in a symmetry-to-reality inference, so what follows is no objection to Healey.

equivalence, and then define a symmetry to be a transformation that (in addition to preserving the laws) preserves them.

Of course, we must not fall into the trap of identifying some privileged physical feature (such as relative position) and say that structures are observationally equivalent when they agree on them. This is what the ontic definitions did, and they fell to the problem of inferential circularity.

Did Ismael and van Fraassen ([2003]) propose an epistemic definition? They defined symmetries to be transformations that (i) preserve the laws, and (ii) 'preserve all *qualitative features* of every model' (p. 379, my emphasis).[35] It depends on what they mean by a qualitative feature. They say that 'qualitative aspects of a physical situation correspond in our terminology to parameters which characterize that situation, *and are directly accessible to us through perception*' (p. 375, my emphasis). This looks like an attempt to delineate some physical features that are directly accessible through perception. If so, the approach is ontic and falls afoul of the problem of inferential circularity. For we typically come to know which physical features are directly accessible by engaging in symmetry-to-reality reasoning (150-years ago, we might have said that spatial distance was directly accessible, but of course today we would not).

Instead, observational equivalence must be defined in epistemic terms that do not depend on the underlying metaphysics, such that we can be in a position to know whether two structures are observationally equivalent prior to knowing anything (via symmetry-to-reality reasoning) about the metaphysics of our world.

Some precedent for this approach might be found in (Galileo [1967]). After describing events in a stationary ship, he asks us to put the ship into smooth motion and says, 'You will discover not the least change in all the effects named, nor *could you tell* from any of them whether the ship was moving or standing still' ([1967], p. 187, my emphasis). The focus here is not that some privileged physical feature (for example, relative position) is preserved, but rather that we cannot tell the difference between the two states. This is the kind of idea that an epistemic definition takes seriously.[36]

But how exactly should we understand this idea that we cannot tell the difference between the two systems? This depends on delicate issues in epistemology and the philosophy of perception. Here is not the place to decide

[35] I should emphasize that they did not present this as a definition of 'symmetry'; their official definition of the term was purely formal. But this is the concept that they claim is at work in the symmetry-to-reality inference, so it is their definition of 'symmetry' as I am understanding the term.

[36] Galileo was discussing two sub-systems, not two complete models or possible worlds. Still, his notion of observational equivalence looks (on the face of it) suitably epistemic.

those issues, but let me roughly outline two approaches to illustrate what an epistemic definition might look like.

6.2 How things look

Observationally equivalent situations look and feel and taste and smell and sound exactly the same.

Suppose we take this seriously. Suppose that each physical system has a 'way it looks' (where 'looks' is now used to express how things are presented to us across all sensory modalities). Then we might say that a transformation $<T, t>$ is a symmetry of L if and only if (i) T preserves L, and (ii) T preserves the way things look. This is an epistemic definition.

This fits well with our paradigm examples of symmetry—rigid translations, uniform boosts, and so on—which all map situations to situations that look the same. If your pasta sauce tastes garlicky, it would still taste garlicky if everything had been shifted three feet over to the right. Moreover, we know this independently of knowing the underlying metaphysics of the world, for example, whether we live in a Newtonian or a Minkowski world. So this definition arguably avoids the problem of inferential circularity.

This basic idea can be developed in a number of ways. One question is whether we should work with an absolute or relative notion of how things look. Things look different from different perspectives, suggesting that we should work with the notion of how things look relative to a given perspective, where a perspective is (perhaps) a location and an angle. But one might, in principle, try working with an absolute notion of how a given system looks—that is, how it looks from nowhere.

Suppose we work with the relative notion of how things look. Then what does it mean to preserve the way that things look from a given perspective? Consider a uniform shift of all matter three feet to the right. Given a perspective P, things will not look the same from P if everything were shifted over since everything would look further over to the right! For everything to look the same, we must also shift the perspective three feet over to the right. In general, then, when asking whether a given transformation $<T, t>$ preserves the way things look, we must transform the perspective by t. (Here, then, is another reason to work with generated, rather than bare, transformations.) It may be that this idea of transforming the perspective by t is not well-defined for some generated transformations (the transformation F discussed in Section 5.2 may be an example). But in that case we can stipulate that a generated function only counts as a symmetry if modifying a perspective by t is well-defined.

This is perhaps the most obvious way of developing an epistemic definition, but it raises a number of questions. If I occupy a perspective then (given my poor eyes) things will look a good deal more blurry than if my eldest daughter occupied it. So perhaps we should work with the doubly relativized notion of how things look from a given perspective to a given person. But then relative to whom should our notion of symmetry be defined? Second, it is not clear that there is a fact of the matter as to how things look (to anyone) from various perspectives. Most perspectives are unoccupied, so the notion is presumably (something like) how things would look to someone were they to occupy it. But consider a very high energy system that a human body could not survive. Is there a fact of the matter as to how things would look to me in that situation?

Perhaps these questions have reasonable answers. But let me outline another approach that might be less controversial.

6.3 Observation sentences

This other approach uses the venerable notion of an observation sentence. Suppose we run a trolley down a slope and measure its progress. We might record the result of the experiment by writing

(1) The trolley came to a halt twenty seconds after being released.

(2) The trolley travelled two metres.

In recording these results we do not mean to commit ourselves to any particular view about the underlying metaphysics of space-time. Regardless of whether we live in a Newtonian space, or a Galilean space-time, or a Minkowski space-time, we want to leave open that these observation reports are correct.

So these sentences must have correctness-conditions that place few constraints on the underlying metaphysics. If we live in a Newtonian space, (2) will be correct if and only if the spatial distance between the start point and the end point of the trolley's trajectory is two metres. If we live in a Minkowski space-time world, (2) will be correct relative to an inertial frame X if and only if the trolley travelled two metres relative to X.[37] And so on.

Are these correctness-conditions truth-conditions? That depends on broader issues in semantics. If one says that they are, then an utterance of (2) can be true even if uttered in a Minkowski world, though it would admittedly not be a perspicuous description of reality. But one might insist that (2) wears its truth-conditions on its sleeves. On this view, (2) is true if and only if two points are a certain spatial distance apart, and so is false if there are no such distances (as is the case in a Minkowski space-time).[38] Still, there is no

[37] This bears unpacking, so that it is clear what it is for x to be two meters from y relative to a frame. But the details do not matter for our purposes.

[38] Boghossian ([2006]) defends this approach to the semantics of location and motion in a relativistic world.

need to choose between these options. For even if one insists that utterances of (2) are false in a Minkowski world, one can still distinguish the good utterances of (2) that satisfy the above correctness conditions and the bad ones that do not, which is all we need.

Now suppose one takes observation sentences like these to constitute our data. Then (roughly speaking) two structures in which the same observation sentences are correct are observationally equivalent. And so we might define a symmetry to be a transformation that preserves the laws and also preserves which observation sentences are correct.

The idea needs a little tweaking. Whether an utterance of (2) is correct when uttered in a Minkowski world depends on the frame of reference in which the sentence is uttered. So instead of asking which observation sentences are correct in a given structure, we should ask which observation sentences are correct in a given structure relative to a given index (for example, a frame of reference or what have you). So, what the correctness-conditions determine is not a mapping from structures to sentences that are correct in that structure, but rather a mapping m from a structure s and an index i to a set of sentences $m(s, i)$ that are correct in that structure relative to that index.

The indices are functioning like perspectives from the last proposal. So when we say that a symmetry must preserve which observation sentences are correct, we mean that given any structure, s, and index, i, the result of transforming s and also transforming i similarly must result in no change about which observation sentences are correct. More precisely, then, let us say that a transformation $<T, t>$ preserves the observation sentences if and only if for any structure, s, and any index, i, $m(s, i) = m(T(s), t(i))$. (Here, yet again, we see the virtue of working with generated transformations.)[39]

So we might then say that a transformation $<T, t>$ is a symmetry of L if and only if (i) T preserves L, and (ii) $<T, t>$ preserves the observation sentences.

Arguably, this epistemic definition avoids the problem of inferential circularity. In order to work out whether a transformation is a symmetry of a law, one does not need to first work out (with symmetry-to-reality reasoning) what the underlying metaphysics of our world is—for example, whether we live in a Newtonian, or Galilean, or Minkowskian world. Instead, the idea would be that one's grasp of observation sentences like (1) and (2) allows one to determine their correctness-conditions in each kind of world, including what kinds of indices correctness is relativized to. These correctness-conditions determine the function m, and with m in hand it is then an *a priori* matter whether the transformation is a symmetry of the law.

[39] As before, it may not be the case that $t(i)$ is well-defined for all transformations.

Moreover this definition arguably avoids the problems that faced the earlier epistemic definition in terms of how things look. Whether (1) and (2) are correct (relative to a given index) does not depend on which subject occupies the perspective of that index or whether she needs glasses. And the correctness conditions of sentences like (1) and (2) may be well-defined even in situations in which there is no fact of the matter about how things would look to a given subject (for example, situations in which I would not survive).

6.4 Observational equivalence

Still, the definition uses the notion of an observation sentence, a notion that some are loath to recognize. But this is not the place to mount a full defense of the notion of an observation sentence.

Indeed, this is not the place to settle on a preferred analysis of observational equivalence. For that clearly involves deep issues in epistemology and the philosophy of perception—issues that would take us too far afield right now. The definitions just presented are merely an approximation of what a completed epistemic definition of 'symmetry' might look like. For now, the conclusion is that the legitimacy of symmetry-to-reality reasoning depends on the existence of some notion of observational equivalence along these lines.

7 Symmetry as an Epistemic Notion (Twice Over)

Suppose for the sake of argument that some epistemic definition like this can be made out. The resulting epistemic reconstruction of symmetry-to-reality reasoning would go like this: We start by asking whether a given determinable property (for example, absolute velocity) is real. To answer this, we ask whether it varies under the symmetries of the laws. Suppose it does. Then what this means (on an epistemic definition) is that there is a systematic way of altering its determinate values such that, given any physically possible system, the result of altering the determinate values like that results in an observationally equivalent system in which the very same laws obtain. It then follows (by the reasoning outlined in Section 4) that the feature is undetectable: there is no physically possible process by which we might discover which determinate values are actually instantiated. So, if there is an alternative theory that does without the feature and that scores well enough on other theoretical virtues, the Occamist norm tells us to believe it.

Let me end by defending this approach from objections and drawing out some consequences.

7.1 Observational equivalence and metaphysics

One might object to using a notion like observational equivalence in a metaphysical investigation into which putative features of the world are real. For observation equivalence is a relative matter: two physical systems that are observationally equivalent for us may be observationally inequivalent for subjects with more refined perceptual capacities. Yet surely whether a putative physical feature is real does not depend on contingencies regarding the build of our eyes! So (the thought is) observational equivalence is unfit for use in metaphysics.

In response, this objection misses its mark. It is correct that which physical features are real does not depend on our perceptual faculties, but the epistemic approach does not deny this truism. The epistemic approach does imply that which physical features we should believe are real depends on which systems are observationally equivalent, and so depends in turn on our perceptual faculties. But this is correct. In theorizing, we can do no better than use the tools we have. Other creatures with better tools may be led to better theories, but this is hardly relevant to what we should believe.

7.2 The Occamist norm revisted

My epistemic reconstruction of symmetry-to-reality reasoning appeals to an Occamist norm. I discussed this norm earlier in Section 4, but what I said there needs some qualification. To see this—and to further understand the mechanics of the epistemic approach—consider the following case of David Baker's (personal communication): Imagine that our best theory says that there are two kinds of particles, Aons and Bons. They behave identically except under very specific conditions C, so we can detect whether something is an Aon or a Bon only by creating those conditions C. Now consider a transformation that changes all Aons that never enter conditions C into Bons and all Bons that never enter conditions C into Aons. By construction this transformation preserves the laws (since Aons and Bons behave identically outside conditions C) and maps models to observationally equivalent models. So it is a symmetry of the laws, according to the epistemic definition. What metaphysical conclusion does our symmetry-to-reality reasoning then recommend?

If our method is to identify models related by a symmetry transformation (or, more precisely, say that they represent the same physical system) then we should conclude that there is no difference between a system in which an Aon never enters conditions C and a system in which a Bon never enters conditions C. This would be an unwelcome result.[40]

[40] It is not uncommon to see this method—of identifying models related by a symmetry transformation—in the literature. For example it is the method employed by (Ismael and van

But this is not the method of the epistemic approach defended here. According to the epistemic approach, we must first identify which features vary under the symmetries, and only then can we conclude that they are undetectable and thus (all else being equal) unreal. So, which features vary under Baker's symmetry?

Not the feature of being an Aon, or of being a Bon. To see this, consider the determinable that consists of the following determinates: {being an Aon, being a Bon}. And let t be the non-trivial bijection on this set. Then Baker's symmetry is not $<T, t>$, where T is the function on structures induced by t. Indeed, $<T, t>$ is not a symmetry, because it switches Aons and Bons that are in conditions C, and hence does not preserve the laws (the laws state that Aons and Bons behave differently in conditions C). So this determinable does not vary under the symmetries of the laws, so our symmetry-to-reality reasoning cannot be used to establish that whether a particle is an Aon or a Bon is undetectable, or that these properties are unreal.

To find Baker's symmetry, we need to take the following determinable D: {being an Aon not in C, being a Bon not in C}. If we let the non-trivial bijection on this set be t^*, then Baker's symmetry is $<T^*, t^*>$, where T^* is the function on structures induced by t^*. So D is variant under the symmetries of the laws, and so (by the reasoning in Section 4) we can conclude that it is undetectable. But this is the right result, since D is in effect defined to be undetectable: it is the determinable property of being an Aon or a Bon in conditions where one cannot detect which is which!

Still, we do not want to say that D or its determinates are unreal. Does our Occamist norm imply that we must say this? Not exactly; it just says (as I put it earlier) that positing undetectable structure is an epistemic vice, in the sense that, all else being (near enough) equal, we should prefer theories that do not posit such structure. So the norm does imply that it is an epistemic vice to posit D, but it says only that we should think that D is unreal if we have an alternative theory that does not posit D and is (near enough) equal on other epistemic virtues.

But it might be objected that it is no epistemic vice at all to posit D. After all, given (almost) any theory, we can cook up a determinable that is defined to be undetectable in much the same way that D is. So unless we want to accuse every theory of having this epistemic vice of positing undetectable structure, we need to qualify our Occamist norm. Why then might it be a vice to posit

Fraassen [2003], p. 80), when they appear to always endorse identifying models related by a symmetry transformation (in my sense of the term 'symmetry'). So Baker's case would seem to be a problem for their view. They discuss a similar case on pp. 384–5 and respond by appealing to a principle of recombination that would prevent identifying the models. It is not obvious how this squares with what they say on p. 380, but I will not pursue this exegetical matter here.

undetectable structure like absolute velocity, but not a vice to posit undetectable structure like D? Well, D is defined in terms of more basic properties: being an Aon, being a Bon, and not being in C. And none of these more basic properties are variant or undetectable in Baker's example. In contrast, velocity either is a fundamental property, or else is defined in terms of cross-temporal spatial distance, which is itself variant and undetectable (in NG). Either way, with velocity we have an undetectable property at the bottom level, whereas with D we get an undetectable property only by stitching together detectable properties.

Our Occamist norm should be sensitive to this. To this end, call a property fundamentally undetectable if and only if (i) it is undetectable, and (ii) it is not defined in terms of detectable properties. Then, instead of saying that it is a vice to posit undetectable structure, our Occamist norm should say that it is a vice to posit fundamentally undetectable structure. Indeed, this revised norm is what we should have endorsed all along. For my argument for the Occamist norm was (very roughly) that discovering that something is undetectable is discovering that we have no empirical reason to think that it is real. But on reflection, this is true only of fundamentally undetectable structure. We have plenty of empirical reason to think that (merely) undetectable features like D are real, for they are defined in terms of properties that are detectable, and hence that we have empirical reason to think are real.

7.3 Consequences

The epistemic approach has a number of consequences regarding the use of symmetries. Let me mention just two.

One consequence is that symmetry-to-reality reasoning is more involved than one might have thought (or hoped). To see this, suppose that 'symmetry' were given a formal or a *de re* ontic definition. Then the question of what the symmetries of a given law are can be settled by straightforward mathematical analysis. This gives the symmetry-to-reality inference a particularly analytic gloss: we do a little mathematical analysis and—hey presto!—we get metaphysical conclusions. It is perhaps because of this veneer of mathematical simplicity that symmetry-to-reality inferences are so popular and yet are given so little justification. But on our epistemic approach, things are not so easy. To work out what the symmetries of a law are, we need to work out which structures are observationally equivalent. And this cannot be settled by mathematical analysis.

To be sure, some cases are uncontroversial: it is clear that classical systems related by a boost or a spatial translation are observationally equivalent, and that general relativistic systems related by a diffeomorphism are too. There is, of course, the question of what justifies this confidence. Perhaps in the case of classical boosts and shifts, the justification is partly empirical, based on our experiences inside trains. Still, no one is likely to question the premise that the situations are observationally equivalent.[41]

But other cases are more controversial. Consider the operation of uniformly multiplying everything's mass (perhaps along with a corresponding change to the gravitational constant or to spatial distances). Suppose we convince ourselves that it preserves the laws of NG. To argue that it is a symmetry, we would need to argue that structures related by such an operation are observationally equivalent. And this is not entirely clear. We certainly do not have the same kind of empirical evidence that we have in the case of uniform boosts—none of us have experienced local environments that differ only in a uniform multiplication of mass. A reason to think that the situations would look the same would have to come from elsewhere, perhaps from a functionalist account of consciousness.

Thus on our epistemic approach, symmetry-to-reality reasoning involves not just mathematical analysis, but also considerations that reach into the philosophy of perception and mind. A far cry from the purely mathematical gloss it is often given!

A second consequence of the epistemic approach is that it dissolves various questions that might otherwise look substantive. To take one example, Roberts ([2008]) discusses the inference from a feature's varying under the symmetries of the laws to the conclusion that it is undetectable (this is the move from (2) to (3) of the symmetry-to-reality inference). He agrees that this is a reliable inference, but he thinks that this is a remarkable fact in need of explanation. The central question of his article is what explains it.

Is the reliability of the inference remarkable? It is if one insists that symmetry is a purely formal or ontic notion. For, in that case, the inference draws a connection between the formalism of the laws and what we can detect. And such a connection would be remarkable indeed! Since Roberts works with Earman's ontic definition of symmetry, it is unsurprising that he considers the connection remarkable.

But, on my view, the connection is hardly remarkable at all. The notion of symmetry in play is, in part, an epistemic notion, with the notion of observational equivalence built into it. As I argued in Section 5, if one does not work with this notion of 'symmetry', the inference is not reliable in the first place. So the reliable connection between symmetry and detection is not a connection

[41] Brading and Brown ([2004]), Kosso ([2000]), and Healey ([2009]) all argue that local symmetries (like diffeomorphisms) are not directly observable. If they are right, the justification for thinking that global operations like boosts are symmetries may be different from the justification for local ones. But this is tangential to the point here, which is just that wherever the justification comes from, we have much justification in the case of boosts and diffeomorphisms to think that they are symmetries.

between mere formalism and detection, but a connection between observational equivalence and detection, which is hardly remarkable. At least, it is not in need of the kind of explanation that Roberts gives it.

8 Conclusion

I have argued that symmetry is an epistemic notion in two respects. First, the symmetry-to-reality inference proceeds by first drawing an epistemic lemma. Second, for the inference to work, the concept of symmetry must be defined partly in epistemic terms.

I focused on the very simple case of boosts and shifts in NG. Even in this sanitized setting, I argued that symmetry-to-reality reasoning is epistemic in these two respects. Thus, when investigating this inference in the context of more sophisticated theories, I suggest that we take this epistemic approach as our starting point.

This may reveal that the prospects for symmetry-to-reality reasoning are better than a reading of (Belot [2013]) might suggest. For Belot starts out with various ontic definitions of 'symmetry' and shows (correctly, I think) that in a number of cases they produce the wrong results. And he concludes that 'it appears that the sort of constraint that knowledge of the symmetries of a theory places on the range of reasonable interpretations of that theory may well be more modest than one might have hoped' ([2013], p. 334). Now, if ontic definitions were in play in the paradigm, sanitized cases of boosts in NG, then Belot would have established the pessimistic claim that this sanitized reasoning does not generalize to more sophisticated settings. But if I am right, then even in those sanitized settings the notion of symmetry at work is epistemic: there is simply no paradigm case of a symmetry-to-reality inference in which a purely formal or ontic notion of symmetry is at work. In which case, the reasoning in these sanitized settings may generalize after all. (Though if what Belot meant by the above quotation is that symmetry-to-reality reasoning is messier than one might have hoped, I agree entirely.)

Acknowledgements

This article went through several incarnations from 2007 to 2013. I doubt I can remember everyone who helped me along the way. I do know that David Baker and Jill North gave me invaluable and detailed comments on earlier drafts, without which this article would probably not exist. My deepest gratitude to them both. Thanks also to Gordon Belot, Adam Caulton, Marco Dees, Kit Fine, Richard Healey, and Robbie Hirsch for their feedback. Some of this material was presented in September 2012 at the 'Structuralism in Physics and Mathematics' conference at the University of Bristol—thanks

to the audience there for their comments. To those I have forgotten, my sincere apologies.

Princeton,
New Jersey, 08544
USA
shamik@shamik.net

References

Baker, D. [2010]: 'Symmetry and the Metaphysics of Physics', *Philosophy Compass*, **5**, pp. 1157–66.

Bell, J. S. [1993]: *Speakable and Unspeakable in Quantum Mechanics*, Cambridge: Cambridge University Press.

Belot, G. [2013]: 'Symmetry and Equivalence', in R. Batterman (*ed.*), *The Oxford Handbook of Philosophy of Physics*, Oxford: Oxford University Press.

Boghossian, P. [2006]: 'What is Relativism?', in P. Greenough and M. Lynch (*eds*), *Truth and Realism*, Oxford: Oxford University Press, pp. 13–37.

Brading, K. and Brown, H. [2004]: 'Are Gauge Symmetry Transformations Observable?', *The British Journal for the Philosophy of Science*, **55**, pp. 645–65.

Brading, K. and Castellani, E. [2007]: 'Symmetries and Invariances in Classical Physics', in J. Butterfield and J. Earman (*eds*), *Philosophy of Physics*, Part B, Amsterdam: Elservier, pp. 1331–67.

Brown, H. and Sypel, R. [1995]: 'On the Meaning of the Relativity Principle and Other Symmetries', *International Studies in the Philosophy of Science*, **9**, pp. 235–53.

Dasgupta, S. [2011]: 'The Bare Necessities', *Philosophical Perspectives*, **25**, pp. 115–60.

Dasgupta, S. [2013]: 'Absolutism vs. Comparativism about Quantity', in K. Bennett and D. W. Zimmerman (*eds*), *Oxford Studies in Metaphysics*, Volume 8, Oxford: Oxford University Press.

Daston, L. and Galison, P. [2007]: *Objectivity*, New York: Zone Books.

Debs, T. A. and Redhead, M. L. G. [2007]: *Objectivity, Invariance, and Convention: Symmetry in Physical Science*, Cambridge, MA: Harvard University Press.

Earman, J. [1989]: *World Enough and Space-Time: Absolute versus Relational Theories of Space and Time*, Cambridge, MA: MIT Press.

Earman, J. and Norton, J. [1987]: 'What Price Substantivalism? The Hole Story', *The British Journal for the Philosophy of Science*, **38**, pp. 515–25.

Einstein, A. [1916]: 'The Foundation of the General Theory of Relativity', in W. Perrett and G. B. Jeffrey (*eds*), *The Principle of Relativity*, London: Methuen.

Feynman, R. [1963]: *Lectures on Physics*, Volume 1, Reading, MA: Addison-Wesley.

Galileo, G. [1967]: *Dialogue Concerning the Two Chief World Systems*, University of California Press.

Greaves, H. and Wallace, D. [2014]: 'Empirical Consequences of Symmetries', *The British Journal for the Philosophy of Science*, **65**, pp. 59–89.

Healey, R. [2009]: 'Perfect Symmetries', *The British Journal for the Philosophy of Science*, **60**, pp. 697–720.

Hoefer, C. [1996]: 'The Metaphysics of Space-Time Substantivalism', *The Journal of Philosophy*, **93**, pp. 5–27.

Howard, D. [1999]: 'Point Coincidences and Pointer Coincidences: Einstein on the Invariant Content of Space-Time Theories', in H. Goenner, J. Renn, J. Ritter and T. Sauer (*eds*), *The Expanding Worlds of General Relativity*, New York: Springer, pp. 463–500.

Ismael, J. and van Fraassen, B. [2003]: 'Symmetry as a Guide to Superfluous Theoretical Structure', in K. Brading and E. Castellani (*eds*), *Symmetries in Physics: Philosophical Reflections*, Cambridge: Cambridge University Press, pp. 371–92.

Kosso, P. [2000]: 'The Empirical Status of Symmetries in Physics', *The British Journal for the Philosophy of Science*, **51**, pp. 81–98.

Kosso, P. [2003]: 'Symmetry, Objectivity, and Design', in K. Brading and E. Castellani (*eds*), *Symmetries in Physics: Philosophical Reflections*, Cambridge: Cambridge University Press, pp, 413–24.

Ladyman, J. and Ross, D. [2007]: *Every Thing Must Go: Metaphysics Naturalized*, Oxford: Oxford University Press.

Lange, M. [2007]: 'Laws and Meta-laws of Nature: Conservation Laws and Symmetries', *Studies in History and Philosophy of Modern Physics*, **38**, pp. 457–81.

North, J. [2009]: 'The "Structure" of Physics: A Case Study', *The Journal of Philosophy*, **106**, pp. 57–88.

Nozick, R. [2001]: *Invariances: The Structure of the Objective World*, Cambridge, MA: Harvard University Press.

Pooley, O. [2003]: 'Handedness, Parity Violation, and the Reality of Space', in K. Brading and E. Castellani (*eds*), *Symmetries in Physics: Philosophical Reflections*, Cambridge: Cambridge University Press, pp. 250–80.

Roberts, J. [2008]: 'A Puzzle about Laws, Symmetries, and Measurability', *The British Journal for the Philosophy of Science*, **59**, pp. 143–68.

Ryckman, T. [2005]: *The Reign of Relativity: Philosophy in Physics 1915-1925*, Oxford: Oxford University Press.

Schaffer, J. [2010]: 'Monism: The Priority of the Whole', *Philosophical Review*, **119**, pp. 31–76.

Sklar, L. [1977]: *Space, Time, and Spacetime*, Berkeley: University of California Press.

Weyl, H. [1952]: *Symmetry*, Princeton, NJ: Princeton University Press.

Wigner, E. P. [1967]: *Symmetries and Reflections*, Bloomington, IN: Indiana University Press.