39

SYMMETRY AND SUPERFLUOUS STRUCTURE: A METAPHYSICAL OVERVIEW

Shamik Dasgupta

39.1 The Method of Symmetry

Symmetry plays a number of central roles in modern physics. As the physicist Paul Anderson famously remarked, "it is only slightly overstating the case to say that physics is the study of symmetry" (1972, p. 394). Here I discuss just one role of symmetry: its use as a guide to superfluous structure, with a particular eye on its application to metaphysics.

What is symmetry? Generally speaking, a symmetry is an operation that leaves its object unchanged in a certain respect. Rotate a square piece of paper 90 degrees in its own plane and it will have the same extension through space; in this sense, the rotation is a symmetry of its extension. But we will focus on the symmetries of *physical theories*, not extensions. Roughly speaking, these are operations on entire physical systems that leave some aspects of the theory unchanged. Which aspects? That depends. Different symmetries preserve different aspects, but the important classes are those that leave the dynamical laws of the theory unchanged. These are known as *dynamical symmetries*.

For example, consider a theory that describes point particles with mass moving through Newtonian spacetime. Suppose the dynamical laws describing their motions are Newton's laws of motion and the inverse-square gravitational force law. In Newtonian spacetime, there is, in addition to a particle's velocity relative to another particle, its velocity *through space* – a quantity that reflects how far it moved (in a direction) over an interval of time independently of its motion relative to other particles.[1] Call this its *absolute* velocity. Now, imagine an operation that takes any possible world consisting of particles in Newtonian spacetime and delivers another world just like it except that everything is uniformly boosted by some velocity V, say 1 mph in a certain direction. At any given time in the boosted world, a particle's absolute velocity is equal to its absolute velocity in the original world at that time plus V. This operation – known as a uniform boost – leaves Newton's laws unchanged: if the original world satisfies those laws, the boosted world does too. (To verify this, note that those laws relate the quantities of mass, ~~the~~ distance at a time, and acceleration; and all these quantities are preserved by a uniform boost.) Hence uniform boosts are dynamical symmetries of this theory.

In truth, this is too quick: as we will see, not *every* operation that leaves the laws unchanged counts a "dynamical symmetry." What other conditions must the operation satisfy? There is little consensus on how the term is to be defined, but we need not settle this here, for a uniform boost is a paradigm example of a dynamical symmetry of this theory, in the sense that every definition of "dynamical symmetry" counts it as such. Other paradigm examples include uniform translations through space (moving everything over in the same direction by the same distance) and uniform translations through time. Our discussion of dynamical symmetries can proceed for now on the basis of these paradigm cases.[2]

I said that symmetry is used as a guide to a superfluous structure. How? Call those features that are unchanged by the dynamical symmetries of a theory the *invariant* features of the theory. Other features are *variant* features of the theory. A uniform boost changes a particle's absolute velocity; hence absolute velocity is a variant quantity of the Newtonian theory above. By contrast, the invariant quantities include acceleration, relative velocity, and spatial distance at a time: all these quantities are unchanged by the dynamical symmetries of the theory. Then the core idea is that *the variant features of a theory are superfluous to it*.

What does this mean? For the purposes of mathematical physics, perhaps just this: that if some part of the mathematical formalism of a theory encodes a variant feature, it can be regarded as physically insignificant and ignored. But for the purposes of metaphysics, the idea is that variance is a sign of unreality – a reason to think that the feature is unreal. Suppose we believed the little theory above and then realized that absolute velocity is one of its variant quantities. The idea is that this would be a reason to think that there is no such thing as absolute velocity after all. To draw this conclusion, one would have to give up the idea that particles move around in Newtonian spacetime and look for an alternative theory of the structure of spacetime on which there is no quantity of absolute velocity. In this way, symmetry is used as a guide to the nature of spacetime.[3]

Metaphysicians should find this intriguing. For metaphysics is in deep need of method, and what we appear to have here is a "method of symmetry": take your best physics, find its dynamical symmetries and hence its variant features, and conclude that those features aren't real.[4] If this method can teach us about the nature of spacetime, we might hope it could generalize to other putative aspects of reality too. If a metaphysician proposes a theory on which reality contains something X, and X can be shown to vary under a dynamical symmetry, the method of symmetry would have us conclude that there is no such thing as X after all.

But how exactly does the method work? And to what other kinds of cases can it be applied? This chapter will survey various approaches to the method of symmetry.

To this end, I focus on the case study of absolute velocity in classical physics. Why trouble ourselves with this tired and fictional example? Wouldn't we be better off drawing metaphysical lessons from a more cutting-edge physics? Yes, ultimately. But if our aim is to understand the *method*, this sanitized example is useful insofar as it abstracts from distracting complications. This is no worse (and, perhaps, no better) than asking whether a white sneaker confirms the hypothesis that all ravens are black. In both cases there may be little interest in the example *per se*; the point is rather to illuminate the logic behind the method.

39.2 Justifications

Let's start by asking what justifies the method of symmetry. Suppose we discover that something is a variant feature of our best physics. Why is this a reason to think that it's not real? What is it about variant quantities that distinguishes them from invariant quantities in this regard? The literature contains at least three answers: that variant quantities are *not objective*, that they are *physically redundant*, and that they are *undetectable*. I'll discuss each in turn.

The idea that invariant quantities are objective has long been associated with symmetry. As Weyl put it, "objectivity means invariance" (1952, chapter 5). This idea was revived by Nozick, who proposed that "an objective fact is invariant under various transformations" (2001, p. 76). So, perhaps the reason to dispense with variant quantities is that they aren't objective.

What does it mean to say that something is "objective"? If it just means that it's real, then it's trivial that quantities that aren't objective aren't real. But then the question is why we should think that variant quantities aren't objective in this sense. That's just the question we started with; no progress has been made.

Another notion of objectivity is perspective independence. To say that a quantity is objective in this sense is to say that it has the same value from all perspectives so that all observers would agree in their measurements of it. Daston and Gallison (2007) argue that this concept played a central role in the thinking of Weyl and other early 20th century physicists. But important as it may be, this conception of objectivity cannot be what justifies our method of symmetry. The reason is that being objective in this sense doesn't have much to do with being invariant. True, velocity isn't objective in this sense since observers moving at a constant velocity relative to one another will disagree in their measurements of velocity: from their own perspectives, each will claim to be at rest while the other moves. But by the same token, acceleration isn't objective in this sense either: if two observers are accelerating relative to one another, each will claim to be at rest from their own perspective while the other is accelerating.

It's unclear, then, whether the concept of objectivity helps to justify the method of symmetry. What we need is a definition of "objectivity" that would apply to invariant but not variant quantities, and which doesn't just *mean* being real. I leave it as an open question whether such a conception can be made out.

Turn now to the idea that variant quantities are *physically redundant*. Here again, we must ask: redundant in what sense? One suggestion is that variant quantities are explanatorily redundant in the sense that they make no difference to how a physical system evolves in time.[5] The idea is tempting because *in*variant quantities do make a difference: two systems that differ initially in inter-particle distances would, according to our simple Newtonian theory, evolve differently over time because the gravitational forces would be different. By contrast, the idea is, absolute velocity isn't like that. Consider two possible physical systems related ~~to~~ a uniform velocity boost. At time t_0 they differ in facts about absolute velocity, yet at all subsequent times, they agree on all invariant quantities including relative positions, relative velocities, accelerations, masses, ~~and charges~~; hence, the initial difference in absolute velocity led to no difference later on. But as Sklar (1974, p. 180) pointed out, this is a mistake: at subsequent times the two systems differ regarding the absolute velocities of each particle! Hence absolute velocity does make a difference to how the system evolves over time after all. One might now say that subsequent differences in absolute velocity aren't *real* differences, but of course, that is the very conclusion we are trying to establish.

Another suggestion is that variant quantities are redundant in the sense that their specific values make no difference to whether the dynamical laws hold.[6] It is true that variant quantities are redundant in this sense: the variant quantities are, by definition, those whose values can be altered while preserving the truth of the dynamical laws. What is less clear is why this is a reason to think that such quantities aren't real. Why should we dispense with a quantity, just because its specific values don't make a difference to whether the dynamical laws hold?

To be sure, anyone justifying the method of symmetry will at some point reach bedrock and appeal to some basic epistemic principle such as induction, modus ponens, or what have you. It would then be unfair to ask them to justify induction! But the principle "dispense with quantities that don't make a difference to whether the dynamical laws hold" isn't epistemically basic in this sense. It may be a sound epistemic principle, but our question is why.

One might answer that if a quantity is redundant in this sense, one can formulate an alternative physics that makes no reference to the quantity. The latter, one might then say, is preferable because it uses "less structure." Compare Newtonian with Galilean spacetime: the idea would be that the latter has less structure than the former and should be preferred on that basis.[7] This is some progress insofar as parsimony is, arguably, a basic epistemic principle. The main challenge is to explain the notion of "less structure"; in particular, one must show that "structure," *so understood*, is what the epistemic principle of parsimony asks us to minimize. As Barrett (2015) and others have argued, it is not clear that there is a notion that does the job.[8]

That leaves the idea that the reason to dispense with variant quantities is that they are *undetectable*.[9] The chief virtue of this approach is that the epistemic principle it rests on is relatively basic and uncontroversial: namely, that in theory choice there should be a presumption against positing quantities that are undetectable. Of course, it may be impossible to dispense with an undetectable quantity without sacrificing other theoretical virtues, in which case one's all-things-considered best theory may retain it. But the point remains that a quantity's being undetectable is *a reason* to prefer a theory that does without it, and that is all the method of symmetry promised us. In what follows, then, I will focus on this approach.

39.3 From Symmetry to Undetectability

The main challenge facing this approach is to show that variant quantities are, in fact, undetectable. This is by no means trivial, but let me sketch one approach.

Consider absolute velocity. Why think it's undetectable? Galileo famously pointed out that physical systems related by a boost would look the same. On the inside, an airplane on the tarmac looks and feels exactly like a plane flying smoothly. Of course, things would look different if you looked outside. But if *everything* were uniformly boosted, including the plane and its surroundings, things would look exactly the same even if you looked out the window. To mark this point, let's say that entire possible worlds related by a boost are *observationally equivalent*.

Still, this doesn't imply that absolute velocity is undetectable. If we can't tell the difference between boosted systems with the naked eye, maybe that's just because our eyes aren't sensitive enough. Perhaps we could build a fancy measuring device that is sensitive to whether it's moving uniformly or at rest. But we can rule this out if uniform boosts are dynamical symmetries of the actual laws. The key is that a measuring device is a physical object and is therefore governed by those laws. So, suppose we built a device in the hope that it would display "Rest" on a computer screen if its absolute velocity ~~was~~ zero, and "Moving" otherwise. And suppose we turn it on and it displays "Rest." Can we infer that it's at rest? No! For consider a uniformly boosted world, in which the absolute velocity of the device is different. If uniform boosts are dynamical symmetries, the same laws are true in the boosted world. Hence the boosted world represents how the device would behave, given the laws governing it, if its absolute velocity differed. And yet the boosted system – by construction – is one in which the device displays "Rest" when we turn it on. After all, a uniform boost preserves all the relative positions of particles, and hence if the pixels on a screen display "Rest" then they'll display the same sign in the boosted system too. Hence the device would say "Rest" *whatever* its absolute velocity was – hardly a "measuring device" to be proud of![10]

This isn't to say that we can't measure velocity *relative* to other bodies. A car's speedometer is sensitive to its speed relative to the road in the following sense: given the dynamical laws governing it, what it displays on its screen is a function of its velocity relative to the road. What the argument purports to show is that, thanks to the symmetries of the laws, it's impossible to build a measuring device that's sensitive to *mere* differences in absolute velocity.

This "symmetry argument" is no ordinary skeptical argument. The ordinary skeptic argues that we can't tell whether we are brains in vats or embodied in a world that is more-or-less as it appears. One response is that so long as our minds are reliably connected to the world, our belief that we are embodied may count as knowledge. Another response is that the hypothesis that we are embodied is a better explanation of our perceptions – simpler, more elegant, and so on – and should therefore be favored on abductive grounds. Yet another response is that in ordinary contexts, the standards for knowledge are sufficiently low that we turn out to know we are embodied after all. But none of these are appropriate responses to the symmetry argument. For one thing, the symmetry argument demonstrates that measuring devices, and by extension, our minds, are *not* reliably connected to facts about absolute velocity. And the competing hypotheses – that our velocity is 1 mph in a certain direction, that it is 2 mph in that direction, and so on – are all equally simple and elegant, so nothing can be said abductively in favor of one hypothesis over the others. For the same reason, lowering the standards wouldn't favor any one of these hypotheses over the others. The symmetry argument is therefore much stronger than an ordinary skeptical argument. No surprise that it is taken seriously by physicists such as Feynman (1963, p. 15), who recognize that it is not just a philosopher's puzzle.

The symmetry argument as stated concerns absolute velocity, but does it generalize to show that *any* variant quantity is undetectable? The argument rested on three facts. First, that the same dynamical laws obtain in the boosted world – without this it wouldn't follow that the boosted system represents how our "measuring" device would behave, given the laws, if its absolute velocity differed. But this was guaranteed by the fact that uniform boosts are dynamical symmetries. So this will generalize for any variant quantity: by definition, there will be a symmetry operation that alters its values yet preserves the dynamical laws.

Second, the symmetry argument required that boosted physical systems are observationally equivalent. Now, we cannot expect this to generalize: we cannot expect that two worlds related by *any* dynamical symmetry of the laws will be observationally equivalent. At least, not given what I've said so far: I said that a dynamical symmetry is an operation that preserves the laws, and operations can do that while changing how the world looks: just consider an operation that maps each physical system to the "null" system containing nothing whatsoever, in which the laws are true vacuously. Still, I also said that *merely* preserving the truth of the laws doesn't suffice for being a dynamical symmetry. More conditions must be added. To generalize the symmetry argument, then, one must show that these extra conditions entail that systems related by a dynamical symmetry are observationally equivalent.

Whether this is true then depends on what definition of dynamical symmetry you work with. But insofar as our aim is to develop the method of symmetry, the natural move is to reverse engineer and *stipulate* that an operation counts as a dynamical symmetry only if worlds related by the operation are observationally equivalent. This could involve building the condition of observational equivalence directly into the definition of dynamical symmetry. Or it could involve providing some other condition – for example, that the operation preserves topological structure – which is then shown to imply this condition of observational equivalence. I will not decide between these two approaches but see Dasgupta (2016) for some considerations in favor of the former.

Third, the symmetry argument required that the hypotheses that distinguish the boosted worlds – that the device is moving at 1 mph to the north, that it is moving at 2 mph to the north, and so on – are equally simple, elegant, common-sensical, and so on; more generally, that they score equally well on every theoretical virtue. This is what distinguished the symmetry argument from an ordinary skeptical argument. To mark this, let us call the competing hypotheses "abductively equivalent," and by extension let us call the boosted worlds abductively equivalent too. Again, we cannot expect this to generalize: we cannot expect that two worlds related by *any* dynamical symmetry *as I've defined the term so far* will be abductively equivalent. But again insofar as our aim is to develop the method of symmetry, the natural approach is to reverse engineer and stipulate that an operation counts as a dynamical symmetry only if worlds related by it are abductively equivalent. As before, I leave open whether the definition contains this condition of abductive equivalence explicitly, or whether it contains some other condition from which abductive equivalence follows.

Of course, if you started off with a fixed idea of what "dynamical symmetry" meant, this reverse engineering may strike you as perverse. Given *your* definition, perhaps the argument doesn't generalize. Fine; there is no need to fight over terms. The fact remains that operations on worlds that satisfy our three conditions – they map worlds to worlds that (i) agree on the laws, (ii) are observationally equivalent, and (iii) are abductively equivalent – have the following property: any quantity altered by such a transformation can be shown, by the symmetry argument above, to be undetectable. Whatever we call these operations, metaphysicians in search of method should find them intriguing!

There may of course be other ways of arguing that variant quantities are undetectable. But for now let us run with the argument we have, which uses the engineered notion of symmetry defined by the three conditions above. I will call these dynamical symmetries, though you are free to substitute with another term if you wish.

39.4 Epistemic Possibility

I said that a dynamical symmetry is an operation on *possible worlds*. But possibility comes in many senses: metaphysical, epistemic, conceptual, logical, and others besides. What is the relevant sense of possibility?

This question is easy to overlook if one studies symmetries mathematically. There, one typically defines a mathematical space of *models* or *states* and then defines a symmetry to be a function on that space.[11] Insofar as the space was mathematically well-defined, there is no *mathematical* question of what entities the symmetry operates on. Still, our philosophical question remains. For a model is just a bit of math, and philosophical conclusions concerning *detection* or *reality* do not follow from math without interpretation. So the question remains as to what the models or states represent. The typical answer is that they represent *possible worlds* or *possible physical states*; our question is what the relevant sense of possibility is.

Metaphysicians may naturally reach for the notion of metaphysical possibility. A uniform boost would then be thought of as a function that maps a metaphysically possible world to one that differs only in facts about absolute velocity.

But a better candidate, I suggest, is the notion of epistemic possibility. On this approach, a dynamical symmetry is an operation on epistemically possible worlds in roughly David Chalmers' sense of a (maximally specific) way the world could be for all we know *apriori*.[12] I know that London once hosted the Olympic Games, but I don't know that *apriori*; hence there is an epistemically possible world in which London never hosted the games. Similarly, I know that water is H_2O, but I don't know that *apriori*; hence there is an epistemically possible world in which water is not H_2O.

To see how this works, consider again the case of Newtonian spacetime. Even if Newtonian spacetime isn't real – even if we *know* it isn't real – we don't know this *apriori*. So there are epistemically possible worlds in which spacetime is Newtonian. Since we can't know *apriori* what trajectories bodies make through Newtonian spacetime, there are many such worlds that differ in the trajectories of bodies. Indeed, there are many such worlds that agree on all relative motions and differ only in a uniform boost – after all, we cannot know a body's absolute velocity *apriori*. We can then understand a uniform boost to be a function on worlds like these: a function that takes an epistemically possible world in which there is a well-defined quantity of absolute velocity and maps it to an epistemically possible world in which all the velocities are uniformly transformed.

Epistemic possibilities have two advantages over metaphysical possibilities. First, our method of symmetry uses dynamical symmetries to establish *epistemic* results about what we can or cannot detect, and the framework of epistemic possibility is well-suited for this purpose. Indeed, in this framework, the argument in Section 39.3 that a variant quantity like absolute velocity is undetectable can be seen, more generally, as an argument that it is *unknowable*. Thanks to being variant, there are epistemically possible worlds that differ in its value; hence we cannot know its value *apriori*. What the argument in Section 39.3 then establishes is that we cannot know its value by observation either. The argument was that whatever output a "measuring" device produces, there are multiple epistemic possibilities in which the output is the same but the value of the quantity differs; hence we cannot know its value on the basis of the observed output either. This last step assumes that the multiple possibilities imply ignorance, an assumption that is trivial if the possibilities are epistemic but not if they are metaphysical.

Second, there may not be enough metaphysical possibilities to make sense of dynamical symmetries. Suppose absolute velocity is not real. And suppose that this is a necessary truth, the idea being that ~~some~~ certain structural features of spacetime are necessary. Then there won't be metaphysical possibilities that differ only in a uniform boost, and so there won't be a function on metaphysical possibilities for a uniform boost to *be*. Yet we still want to say that boosts are symmetries; indeed, that may be our reason for thinking that absolute velocity isn't real in the first place! Epistemic possibilities avoid the problem: even if there are no metaphysical possibilities that differ in a uniform boost, there are epistemic possibilities that do.

39.5 Empirical Symmetries

The "method of symmetry," as we now have it, is this. We take our best physics. We look at its laws and find its dynamical symmetries and hence the variant quantities. We then argue that those quantities are unknowable. And we take that as a reason to think that the quantity isn't real.

This last step should now be clearer. There is no claim that being unknowable *logically implies* that the quantity isn't real: after all, there are epistemically possible worlds in which it is real. For the same reason, this is not a verificationist argument that talk of the quantity is meaningless. The idea is just that positing undetectable quantities is a bad-making feature of a theory so that theories that dispense with the quantity are preferable in that respect. Of course, whether the latter theory is all-things-considered preferable then depends on assessing its other virtues and vices, and much argument might be had on that matter. Still, we have on our hands a principled method that yields *reasons* for various metaphysical theses. In metaphysics, that counts for a lot!

But if this is our method, it arguably generalizes to a broader class of operations. To see this, suppose that absolute velocity is *not* a variant quantity. To take a fictional example, suppose that color is a primitive property of bodies, and suppose there's a dynamical law to the effect that things take on a greenish hue when their absolute velocity is zero but not otherwise. Uniform boosts would not then be dynamical symmetries because they don't preserve this law: they map worlds in which the law holds to worlds in which objects take on a greenish hue when they have some non-zero velocity instead. Absolute velocity, then, would be an *invariant* quantity. Does this mean that we could measure whether something is at absolute rest by looking to see whether it is greenish? Unfortunately not. For if we see something greenish, we can infer that its absolute velocity is zero only on the basis of a *theory*, namely that green correlates with zero velocity. Contrast this with the theory that green correlates with some non-zero velocity, say 5 mph in a certain direction. If this theory were true, green would indicate this non-zero velocity instead. But how could we ever tell which theory is true? All we would ever *see* is that the bodies with a greenish hue are all at rest relative to one another, and moving relative to other things. True, this means that if there are absolute velocities, one of them would correlate with green. The trouble is, there'd be no way of telling which one it is.

Or suppose there was a velocity-dependent force law. Suppose it implies, of a particular system, that if its center of mass is at absolute rest then it'll evolve one way, while if the velocity is non-zero it'll evolve differently. And suppose the difference between the two evolutions is directly observable: a different pattern of ink on paper is produced in each case, say. In this case, again, uniform boosts wouldn't be dynamical symmetries, so absolute velocity would be an invariant quantity. Can we then measure absolute velocity by observing how the system evolves? No again, and for the same reason. We can only infer that a system exhibiting the first pattern is at absolute rest on the basis of the *theory* that that pattern correlates with absolute rest. Contrast this with a theory on which that pattern correlates with some non-zero velocity instead. Both theories are consistent with what we'd *see*, namely that all systems that display that first pattern are at rest relative to one another, and moving relative to systems that don't. But there'd be no way of telling what the absolute velocity of those former systems is.

The point is that *measuring* a quantity involves inferring its value on the basis of (i) a measurement outcome, and (ii) the physical laws that generated that outcome.[13] In Section 39.3 we assumed that we knew the laws; the problem was that the same outcome would be produced regardless of what the value of the quantity was. Here the problem is that we can't know what the laws are in the first place.

With these fictional laws, then, absolute velocity would be unknowable even though uniform boosts wouldn't be *dynamical* symmetries. Still, boosts would be symmetries in a broader sense. For one thing, there's an intuitive sense in which boosts preserve the "form" of these laws. In the case of color, all boosted worlds have a law of the form 'green correlates with V'; they just differ on the particular velocity V. Indeed, theorists of symmetry often focus on transformations that preserve the form in something like this sense.[14] But what exactly is form? For our purposes, the important fact is that we can never know which of the various color laws obtains. This requires that each law is abductively equivalent in the earlier sense. Consider the range of color laws: that green correlates with absolute rest, that it correlates with velocity 1 mph in a certain direction, that it correlates with velocity 2 mph in that direction, and so on. The idea behind the argument above is that each law is equally simple, explanatory, and so on, so there is no reason to believe one over the others. (You might object that we should favor the law that green correlates with absolute rest because it's simpler than the rest. But then you'd also have to say that the hypothesis that the center of mass of the universe is at absolute rest is simpler than its boosted alternatives, in which case boosted worlds aren't abductively equivalent and the argument that absolute velocity is unknowable in Section 39.3 doesn't go through. But surely we can't know our absolute velocity on this basis! What this shows is that the sense in which the rest hypothesis is simpler is epistemically insignificant.)

Thus, we can define the *empirical symmetries* of a theory just like dynamical symmetries with the one exception that the transformed world isn't required to have the *same* laws as the original world, just *abductively equivalent* laws. Since any law is abductively equivalent to itself, it follows that any dynamical symmetry is also an empirical symmetry. But the reverse is not the case: uniform boosts are empirical symmetries of the color law, but not dynamical symmetries. So the notion of empirical symmetry is more general, and the idea would be that any quantity that varies under the *empirical* symmetries of the laws is unknowable.

This then gives us a more expansive method of symmetry. We take our best physics, look at its laws and find its *empirical* symmetries, and argue that quantities altered by *those* symmetries are unknowable. As before, we then take this as a reason to think that the quantity isn't real. This method is more expansive insofar as it'll lead us to dispense with more quantities. If this leaves you feeling queasy, be assured that there are limits. Reichenbach's (1958) transformation that maps a non-Euclidean world to a Euclidean world with universal forces, for example, isn't an empirical symmetry. The worlds are observationally equivalent, but not abductively equivalent because they have different abductive virtues (the former is simpler, in one sense of the term, while the latter is closer to pre-theoretic belief about geometry).

Note that this move to empirical symmetries hangs on our decision, in Section 39.2, to think of the method of symmetry as eliminating variant quantities because they're *undetectable*. If we had eliminated them because they're *physically redundant*, this move to empirical symmetries would not make sense. For in the color law above, absolute velocity is *not* redundant in any of the senses surveyed there: it *does* make a difference to whether a body will take on the greenish tinge; its values *do* make a difference to whether the law holds; and so on. I said in Section 39.2 that this justification of the method of symmetry faces problems anyway; my point here is that if one justified the method this way, one would be restricted to using dynamical symmetries. The move to empirical symmetries makes sense only if thinks of the method of symmetry as eliminating *undetectable* quantities, for as we've seen these include those that are variant under the empirical symmetries, not just the dynamical symmetries.

This method of empirical symmetry can lead to substantial revisions in physics. Suppose we started with a theory on which green correlates with absolute rest. Suppose we then use the method of empirical symmetry and dispense with absolute velocity. We must then revise our theory. We can no longer say that green correlates with any particular velocity. Rather, we can only say that things that take on a greenish hue are all at rest relative to one another, and are moving relative to things that don't. This is a very different theory. For one thing, in the original theory, a body's state of motion *determines* whether it takes on a greenish hue, but in the revised theory this is not so. If we take a possible system in which two things are at rest relative to one another, our revised theory doesn't determine whether they'll have the greenish hue; it just determines that either both will or neither will. Thus, the method of empirical symmetry led from a theory that was deterministic to one that is not.

This raises the question of whether the new theory would be all-things-considered better than the old theory. True, the old theory had the vice of containing unknowable structure, but the new theory is indeterministic. Is this a vice? If so, is it worth it? More generally, if the method of empirical symmetry leads us to new physical theories with surprising features, will the new theory be all-things-considered better?

This can only be settled on a case by case basis, based on a close examination of the rival theories. But we can say this: we cannot object to the new theory *on the basis of observation*. In the case of color, we cannot complain that we *observed* bodies behaving in the deterministic manner described by the old theory. For, the method of empirical symmetry shows that we observed no such thing. It shows that we never saw that the greenish hue correlates with *absolute rest* rather than any other velocity, we just saw that things with the hue are at rest relative to one another. If we previously *thought* we had observed that green correlates with absolute rest, that was based on a mistaken characterization of the observations – perhaps we had a deeply ingrained belief that some reference body was in a state of absolute rest so that when we saw something at rest relative to it we mistakenly chalked that up to seeing it at absolute rest. In any event, the point is that an objection to the indeterminism cannot be based on observation; it must be based on considerations of another kind.[15]

More generally, the method of empirical symmetry assures us that the new theory is empirically adequate if the old theory was, for the new theory dispenses only with quantities that are undetectable. Thus, any surprising feature of the new theory cannot be said to be inconsistent with observation.

39.6 Metaphysical Applications

I have outlined two methods of symmetry, one based on dynamical symmetry and the other on empirical symmetry. Which method should we use? This depends on our purposes. Philosophers of physics sometimes ask what the appropriate spacetime structure is *for a particular set of laws*. For example, what is the natural spacetime structure for classical Newtonian mechanics? In this case, dynamical symmetries are more appropriate. We want to look at those worlds *in which the laws obtain* and see what structure we can dispense with. Insofar as the method is to dispense with undetectable structure, we can think of us as asking: conditional on these laws, what would be undetectable? Dynamical symmetries are perfectly suited for this purpose.

But we may alternatively ask what spacetime structure is indicated *by a certain class of observations*. Given the kind of observable phenomena~~,~~ there'd be in a Newtonian world, for example, we may ask what we should believe about the structure of spacetime if we observed those phenomena. In this case, the empirical symmetries are more appropriate. We want to look at the worlds that agree on those observational phenomena and ask what would be undetectable, and in doing so it makes sense to consider worlds with different laws if we couldn't know which of those laws obtained.

To the extent that metaphysicians engage with physics at all, the second question is the more pertinent. They want to know what's real. To this end, the method of symmetry has them ask which features are undetectable. But it would be odd to hear that as asking what would be undetectable conditional on a particular set of laws if we could never know that those laws obtain! Thus, metaphysicians should be particularly interested in the method of empirical symmetries.

What metaphysical questions might this method illuminate? Questions about the structure of spacetime are most familiar. But if we use empirical symmetries, not dynamical symmetries, we may be led to surprising conclusions. Consider classical Newtonian mechanics. And consider a time-dependent rotation: an operation that takes a physical system and puts it into constant rotation over time around some axis. This is not a dynamical symmetry. Indeed, that is the point of Newton's bucket argument: if the water in a bucket is flat, then according to Newton's laws it would go concave if spinning; hence a world that differs only in its state of rotation but not in the concavity of water in buckets would violate Newton's laws. This leads many commentators to the view that the natural spacetime structure for Newtonian mechanics contains a notion of absolute rotation.[16] Still, one might argue that time-dependent rotations are *empirical* symmetries of the theory. As Sklar (1974, p. 230) pointed out, we don't directly *see* that the concave water is rotating; what we see is that it's rotating *relative to water that remains flat*. More generally, we don't directly see which particular state of rotation correlates with inertial effects; we just see that bodies undergoing different inertial effects are rotating relative to one another. Thus there are many possible theories consistent with the data: that no inertial effects correlate with *zero* absolute rotation, that it correlates with a *non-zero* state of rotation R, and so on. *If* these theories are abductively equivalent, then time-dependent rotations are *empirical* symmetries and absolute rotation is unknowable after all. That is a big "if" (to put it mildly!) and I won't try to settle it here; the point is that for someone using the method of empirical symmetry, this is the pertinent question.

The upshot is that if one asks what the natural spacetime structure for Newtonian mechanics *as it is standardly written*, the answer might be that it must contain a notion of absolute rotation. But if one asks what spacetime theory would be best confirmed if one lived in a world with Newtonian observational phenomena, the answer might be the opposite.

I've focused on symmetries that act on spatiotemporal quantities such as boosts and rotations. But the same goes for symmetries that act on other quantities or fields. These include symmetries that act on mass and charge (see Dasgupta, 2013); permutation symmetries in quantum mechanics and statistical mechanics (see entries by Ladyman and Caulton, this volume), and gauge symmetries (see the entry by Teh, this volume). In all these cases we must distinguish the dynamical from the empirical symmetries. Suppose one of these operations is *not* a dynamical symmetry of some theory as standardly written. What that means is that *conditional on the theory* we can detect the quantity altered by the operation. But that leaves open the possibility that the operation is an *empirical* symmetry, in which case we can't detect the quantity because we can't know whether the theory is true.[17]

References

Albert, D.Z. (2015). The difference between the past and the future. In D.Z. Albert (ed.), *After Physics*. Cambridge, MA: Harvard University Press, 31–70.

Anderson, P. W. (1972). More is different. *Science*, 177(4047): 393–396.

Baker, D. (2010). Symmetry and the metaphysics of physics. *Philosophy Compass*, 5: 1157–1166.

Barrett, T. (2015). Spacetime structure. *Studies in History and Philosophy of Modern Physics*, 51: 37–43.

Belot, G. (2013). Symmetry and equivalence. In R. Batterman (ed.), *The Handbook of Philosophy of Physics*. Oxford: Oxford University Press.

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: Elsevier, pp. 1331–1367.

Chalmers, D. (2012). *Constructing the World*. Oxford: Oxford University Press.

Dasgupta, S. (2013). "Absolutism vs Comparativism about Quantity" *Oxford Studies in Metaphysics*, 8 (2013): 105–148.
Dasgupta, S. (2016). Symmetry as an epistemic notion (twice over). *The British Journal for the Philosophy of Science*, 67: 837–878.
Dasgupta, S. (2021). How to be a relationalist. *Oxford Studies in Metaphysics: Volume 12*. Oxford: Oxford University Press, pp. 113–163.
Dasgupta, S. and Turner, J. (2017). Postscript. In E. Barnes (ed.), *Current Controversies in Metaphysics*. New York: Routledge, pp. 35–42.
Earman, J. (1989). *World Enough and Space-Time: Absolute versus Relational Theories of Space and Time*. Cambridge, MA: MIT Press.
Feynman, R. (1963). *Lectures on Physics: Volume 1*. Reading, MA: Addison-Wesley.
Healey, R. (2009). Perfect symmetries. *The British Journal for the Philosophy of Science*, 60: 697–720.
Ismael, J. (this volume). *Symmetry and Superfluous Structure: Lessons from History and Tempered Enthusiasm*.
Ismael, J. and B. van Fraassen (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. 371–392.
Knox, E. (2014). Newtonian spacetime structure in light of the equivalence principle. *The British Journal for the Philosophy of Science*, 65(4): 863–880.
North, J. (2009). The 'structure' of physics: A case study. *The Journal of Philosophy*, 106: 57–88.
Nozick, R. (2001). *Invariances: The Structure of the Objective World*. Cambridge, MA: Harvard University Press.
Reichenbach, H. (1958). *The Philosophy of Space and Time*. Mineola: Dover Publications.
Roberts, J. (2008). A puzzle about laws, symmetries, and measurability. *The British Journal for the Philosophy of Science*, 59: 143–168.
Saunders, S. (2013). Rethinking Newton's *Principia*. *Philosophy of Science*, 80: 22–48.
Sklar, L. (1974). *Space, Time, and Spacetime*. Berkeley: University of California Press.
Wallace, D. (2020). Fundamental and emergent geometry in Newtonian physics. *The British Journal for the Philosophy of Science*, 71(1): 1–32.
Weyl, H. (1952). *Symmetry*. Princeton, NJ: Princeton University Press.

Further Reading from the Editors

H. Weyl, *Symmetry* (Princeton, NJ: Princeton University Press, 1952) gives the view of a great physicist and mathematician. J. Earman, *World Enough and Space-Time: Absolute versus Relational Theories of Space and Time* (Cambridge, MA: MIT Press, 1989) provides an overview of spacetime structure and symmetries. More general thoughts about the connections between symmetry and metaphysics can be found in D. Baker, 'Symmetry and the Metaphysics of Physics' (*Philosophy Compass* 5: 1157–1166, 2010). A longer overview is G. Belot, 'Symmetry and Equivalence' in R. Batterman (ed.), *The Handbook of Philosophy of Physics* (Oxford: Oxford University Press, 2013). For a development of the views expressed in this chapter, see S. Dasgupta, 'Symmetry as an Epistemic Notion (Twice Over)' (*The British Journal for the Philosophy of Science* 67: 837–878, 2016).

Notes

1 For a description of Newtonian spacetime see Samaroo (this volume).
2 For representative examples of definitions in the philosophical literature, see Earman (1989, pp. 45–46) and Belot (2013).
3 See Ismael, this volume, for a description of "Galilean" spacetime in which there is no quantity of absolute velocity.
4 Statements of this "method of symmetry" abound in the literature; for a small sample see Earman (1989, p. 46), Ismael and van Fraassen (2003), Roberts (2008), North (2009), Healey (2009), Baker (2010), Belot (2013), and Dasgupta (2016).
5 See for example Baker (2010).
6 Earman (1989, p. 46) might be read as advancing this idea.
7 See North (2009) and Belot (2013) for explicit statements of this idea. It may also be what Earman (1989, p. 46) had in mind.
8 See also Dasgupta and Turner (2017) for some doubts as to whether "minimizing structure" is the norm driving the move from Newtonian to Galilean spacetime.
9 This idea has been defended explicitly by Ismael and van Fraassen (2003), Roberts (2008), Dasgupta (2016), and Ismael (this volume).

10 This style of argument is sometimes gestured at by physicists such as Feynman (see his 1963, p. 15). But it was spelt out explicitly by Roberts (2008); my presentation here is based on Roberts'.

11 The use of mathematical spaces of models or states is widespread, but for some explicit examples see Earman (1989, p. 45), Brading and Castellani (2007, p. 1342), and Baker (2010, p. 1158).

12 See Chalmers (2012) and references therein for more on this notion of epistemic possibility.

13 See Albert (2015), who emphasizes that one also needs information about the initial state of the measuring device. But I bracket this here for simplicity.

14 See, for example, Brading and Castellani (2007, p. 1342).

15 I discuss this issue of indeterminism in Dasgupta (~~forthcoming~~).

16 See for example Saunders (2013), Knox (2014), and Wallace (2020).

17 I am deeply grateful to Thomas Barrett, Daniel Berntson, and Alastair Wilson for their extremely helpful and timely comments on an earlier draft of this paper.