

Symmetry and the Metaphysics of Physics

David John Baker

Department of Philosophy, University of Michigan

djbaker@umich.edu

June 22, 2010

Abstract

The widely held picture of dynamical symmetry as surplus structure in a physical theory has many metaphysical applications. Here I focus on its relevance to the question of which quantities in a theory represent fundamental natural properties.

1 Introduction

The proper role of physics as evidence for metaphysicians is the subject of much dispute (see e.g. Maudlin (2007) and Ladyman and Ross (2007)). But there is one topic in metaphysics on which almost all will agree that physics is most qualified to comment. We commonly think of our universe as made up (at least partly) of some fundamental natural properties, quantities which each take on one of a range of physically possible values. Typical examples include mass, position, energy and electric charge, though it may be that these will be replaced as fundamental quantities by something like the amplitude of the quantum wavefunction or the configuration of superstrings. Although this conception of the world as made up of natural quantities has been challenged, most famously by Goodman, it remains an important part of most accepted systems of metaphysics.

It is therefore a very good question, according to these systems, what the fundamental natural quantities actually are, and perhaps also what they might be in other physically possible worlds. And it is normally thought that fundamental physics provides our best means of answering this question.

This means it is of paramount importance, metaphysically, to develop a method for extracting from a physical theory a description of the fundamental quantities according to that theory. One promising method, which has received significant attention in the philosophy of physics literature, appeals to the concept of symmetry. My threefold goal here is to introduce this method, provide a *prima facie* argument in its favor, and discuss some cases in which it seems to work rather well.¹

The first step will be to make clear what symmetry signifies in physics. The word is used in a few different ways, but I'll be concerned here with *dynamical symmetries*. This concept admits an intuitive as well as a formal definition; the intuitive definition is: symmetries of a theory are transformations that preserve its laws.

We have a rough sense of what it means for a transformation to preserve an object – that is, map it to itself. A square with four identical corners and four identical sides is preserved in this way by 90-degree rotations, but not by 45-degree rotations, which map it to a diamond instead of an identical square. We say that the 90-degree rotations leave the square invariant by mapping it to a qualitatively identical geometric object. In a similar way, transformations like real-world rotations can leave the laws of physics invariant. Newton's theory accurately predicts how long it will take my pencil to fall if I let go of it, and these predictions are preserved even if the pencil and myself are both rotated by some angle in space. Since all the predictions of Newtonian mechanics are preserved in this way by rotations, the *group* of all rotations is a symmetry of Newtonian physics.²

The formal definition requires a bit of mathematical machinery. Complete physical theories like Newtonian mechanics, relativity and quantum mechanics can be formulated in a mathematical arena called a *state space*. We use that name because every element in state space stands for a physically possible (instantaneous) state of the world according to our theory. The experimental information we get from a theory comes in the form of predictions about how states will change over time. We call this account of temporal change a theory's

¹The method I present here is one of many proposals, and it is perhaps less nuanced than most. I see its simplicity as a virtue, but many others would call it oversimplified. For an opposing view of symmetries, see Maudlin (2002, 1-7).

²By “group” I mean something more interesting than a family or set. A group G is a set with a distinguished “identity” element e and a “product” (two-place function) (\cdot) such that $e \cdot g = g$ for all $g \in G$ and every $g \in G$ has an “inverse” g^{-1} such that $g^{-1} \cdot g = e$. The set of all rotations in three-dimensional space clearly form a group (called $SO(3)$) if we define the product of two rotations as the result of performing one rotation and then the other.

dynamics; mathematically, the dynamics is sometimes represented by time-indexed transformations $U(t')$ on state space that takes a state at time $t = 0$ to the state it will change into at time t' . So a theory's dynamics is a mapping from states to states. Transformations like rotations are also given by mappings T from states to states. Symmetries are then given by transformations that leave the dynamics (diachronic laws) unchanged. Mathematically, this means they must commute with the dynamics, so that $U(t')T = TU(t')$ for every symmetry transformation T and every time t' . It will turn out that transformations like these meet our intuitive condition of leaving the laws invariant.³

With the concept of symmetry in hand, we now move on to the question of why these transformations should have any systematic metaphysical importance. In short, it will turn out that states related by symmetries must share identical values of all fundamental quantities.

2 Using symmetry to find fundamental quantities

What features distinguish fundamental quantities from non-fundamental ones? Most importantly, fundamental quantities are supposed to ground or explain objective similarity between objects and states of affairs, in a way that non-fundamental properties cannot. For example, two electrons are objectively similar because they are both negatively charged (and thus agree on a fundamental quantity), but not because they are both mentioned in this article.

The relation of similarity grounded by fundamental quantities on this picture is supposed to be objective – it's up to Nature, not us, to determine which things are qualitatively similar. One plausible necessary condition for objectivity is invariance under descriptive changes. The relation of similarity between objects shouldn't change depending on which language, coordinate system, etc. I use to describe them.

A coordinate system (at least in familiar Euclidean geometry) is a piece of *surplus structure*. It's something extra we add to the theory to aid us in describing coordinate-independent facts. It has been suggested, for example by Belot (2001), that physical quantities which vary under symmetry transformations are like coordinate systems in this way. In other

³Time-indexed dynamical transformations of this sort are not always well-defined. In general, this article will follow most of the literature in ignoring and oversimplifying a very deep problem: given a physical theory, how do we determine its symmetries? See Belot (in progress) for a sophisticated look at these problems.

words, the symmetries of a physical theory are a guide to surplus structure in the theory's formalism. If this is correct, and if changes in surplus structure are generally (as in geometry) mere descriptive changes, it follows that physical situations related by symmetries must be qualitatively identical. And if this is right, then physical quantities that change under symmetry transformations (i.e., that are not invariant) must not be fundamental quantities. Qualitatively identical objects or worlds cannot disagree about the fundamental quantities.

Why believe that symmetry transformations change only surplus structure? This question is somewhat murky, and my answer will signify a good place for those who disagree with my approach to get off the boat. It's an answer I find compelling nonetheless. It rests on the notion that the language of fundamental physics is complete in a particular sense: in a satisfactory physical theory, the fundamental quantities are all dynamical difference-makers. Difference-makers, that is, in the sense that differences in the fundamental quantities must make some difference in how the state of a (physically possible) world evolves in time.

It would be bizarre if by fixing the values of all fundamental quantities, we could not thereby fix (if not deterministically, at least probabilistically) how a physical world will evolve in time. In my view, the converse seems equally plausible: if objects (or worlds) are fundamentally different, fundamental physics should recognize and explain that fact. A theory in which some fundamental quantities are completely epiphenomenal is, for that reason, an unsatisfactory theory. A quantity is epiphenomenal, I take it, if it can be left out of a complete dynamical explanation of the world's evolution in time. It follows that quantities with no bearing on the evolution of the world's physical state must not be fundamental. Dynamical explanations can make do without these quantities, so their claim to fundamentality is suspect.

The present approach may also be bolstered by an epistemological argument of the sort offered by Healey (2006). Without addressing the metaphysical question of whether possibilities related by symmetries could differ qualitatively, Healey argues that it is never justifiable to accept a theory according to which they do differ. He bases this on a preference for theories which are "uniquely realized," which means that all of their theoretical terms can be defined using the method proposed by Lewis (1970). Theories which are not uniquely realized ("multiply realized" theories) require that we use demonstrations if we want to name theoretical entities or properties with indistinguishable roles in the theory. For example, if water and twater are distinct substances with indistinguishable roles in our theory, we will

have to name them by pointing at particular pools of liquid, rather than defining one (per Lewis's method) as the unique realizer of the water role and the other as the unique realizer of the twater role.

Lewis (1970, 146) claims that "A uniquely realized theory is, other things being equal, certainly more satisfactory than a multiply realized theory," and Healey adamantly agrees. When it is possible to formulate a theory with terms whose reference is uniquely fixed by their role in that theory, Healey thinks we should do so. To do otherwise would be to flout Ockham's dictum to seek out simpler theories; more importantly for Healey, it would be to turn against the scientific realist tenet that the explanatory power of a theory is what gives us reason to believe in its truth. A multiply realized theory is unsatisfactory because it includes "elements purporting to represent real structures but that play no role in contributing to the theory's success," since nothing about the theory's success hinges on how these elements of the theory are defined. To describe the world using a multiply realized theory requires that we define its theoretical terms arbitrarily, in such a way that at least one other definition would do just as well.

By treating possibilities related by symmetries as fundamentally identical, this arbitrariness can be avoided. In this case, every fundamental quantity will possess a unique theoretical role – otherwise it would be related by a symmetry to any other quantities with identical roles. So the present approach can ensure that any theory of the qualitative facts is uniquely realized – and thus epistemically preferable to the alternatives, by Healey's lights.

Other arguments to this same conclusion, that qualitative features are invariant under symmetries, have also been advanced. Roberts (2008) provides a convincing argument (really an explanation of a well-known fact) that measurable quantities must be invariant. So if one holds to a modest verificationism according to which qualitative features of the world must be measurable in principle, the conclusion about symmetries follows. A similar point is made by Ismael and van Fraassen (2003), who explicitly endorse such verificationism. I don't necessarily agree with this premise myself, but those who do should embrace my approach to symmetries.

3 Spacetime symmetries: Leibniz equivalence

The most straightforward – and fruitful – consequences of my preferred approach to symmetries can be found in the interpretation of spacetime theories. A symmetry-based understanding of these theories allows us to separate the question of whether space (or spacetime) exists from related questions about the modal features of its parts (points and regions). One of Leibniz’s most important theses about Newtonian mechanics is really a claim about the modal properties of points, and it can be defended decisively if we treat the symmetries of mechanics as surplus structure. This thesis of Leibniz equivalence can be formulated analogously in relativistic theories of spacetime, and remains plausible there for the same reasons.

A major point of contention between Leibniz and Clarke was the possibility of certain counterfactual arrangements of matter in space. Clarke held that all of the spatial relations between material objects might have been just as they actually are, but the positions of these objects shifted uniformly with respect to space itself. We all might have been located a foot to the left of where we actually are, and everything else been the same. Leibniz denied this possibility.

Leibniz defends his view by appeal to the principle of sufficient reason (PSR). In an important paper, Belot (2001) shows that the same work can be done by a somewhat revisionary version of PSR that treats qualitatively indiscernible worlds as identical. Belot’s PSR can always be defended against purported counterexamples, he claims, if we insist (as in the present approach) that symmetries reveal surplus structure.

In applying his PSR to Newtonian mechanics, Belot presents an important mathematical structure that can be constructed for any physical system with symmetries: the *reduced state space*. Recall that we defined a symmetry as a mapping on state space (the space of “physically possible worlds” of a theory). The reduced state space is the space of equivalence classes of worlds related by symmetries. In other words, it is exactly the structure we get if we identify worlds related by symmetry transformations – qualitatively identical worlds, on the present approach.

How does this work in Newtonian mechanics? As Belot explains, the state space Newton used for his mechanics of N particles is $6N$ -dimensional: there are six degrees of freedom for each particle, including its position and momentum in each of three dimensions in physical

space. (Note that in this case each degree of freedom corresponds to a fundamental quantity.) This state space includes many qualitatively identical states related to each other by global translations or rotations of all the matter in space, or by global increases in the velocity of all matter (which we call Galileian boosts). The set of all these transformations taken together is the Galileian group, the symmetry group of Newtonian mechanics.

In constructing the reduced space of the Newtonian state space by identifying states related by Galileian symmetry transformations, we effectively reduce the theory's degrees of freedom. For example, there are distinct states in Newton's state space that differ only by the position (in absolute space) of the universe's center of mass. All these states are treated as a single state in the reduced space. As a result, three degrees of freedom – the ones needed to specify the position of the center of mass – are left out of the reduced space. In total, ten degrees of freedom can be eliminated in this way, corresponding to the position and linear motion of the center of mass and the rotational motion of the universe as a whole.

Thus while Newton believed his N -particle mechanics required $6N$ fundamental quantities, $6N - 10$ will do perfectly well. This revision will in turn require a picture of space and time according to which quantities like the position and linear motion of the universe's center of mass are not needed to give a full description of a physical possibility. That sort of picture is provided by so-called Galileian spacetime, in which acceleration, but not position or velocity, is absolute.⁴ In this way, the present approach to symmetry points toward a superior picture of space and time compatible with the success of Newton's theory. The key was to identify states that are related by spacetime symmetries as the same physical possibility. The inference from spacetime symmetry to physical equivalence has sometimes been called "Leibniz equivalence," after his early arguments to similar conclusions.

The present approach can be applied similarly, and even more successfully, to our most advanced spacetime theory: general relativity (GR). Unlike previous spacetime theories, GR doesn't occur against the fixed background of a single spacetime. Instead, many distinct spacetimes (with distinct symmetries) are solutions of the same theory. The symmetries of these individual spacetime solutions often provide interesting (although physically contingent) information. But the interesting symmetries for purposes of the present approach are

⁴As Belot points out, we can even do without absolute acceleration if we count global variables for the total angular momentum of the universe as fundamental quantities – although this further parsimony is not entailed by the present approach to symmetry.

the symmetries of GR as a whole: diffeomorphisms.

A diffeomorphism is simply a differentiable map taking all points in a spacetime to all points in another (or the same) spacetime. In effect, a diffeomorphism is a re-labeling or re-assignment of spacetime points which leaves unchanged spacetime's matter content and its metric structure (distance relations). For example, a diffeomorphism must leave unchanged the distance between massive objects, but may alter which spacetime points those objects occupy – in doing so it also alters the distance relations between the points. On the present approach to symmetries, a physical spacetime should be given by a diffeomorphism equivalence class of mathematical spacetimes. The diffeomorphism symmetry of GR thus tells us that it makes no physical difference which points occupy which physical roles, so long as the same physical roles and relations are instantiated.

What would it be like for the assignment of physical roles to points to be of metaphysical significance? This would require, at a minimum, the existence of primitive facts about the trans-world identity of points. For, if diffeomorphisms can relate different possibilities, it must be possible for spacetime points to have their qualitative features reversed (or at least for two different points to have the same qualitative features in two different worlds). This might require a sort of haecceitism about spacetime points; at least it requires the existence of apparently haecceitistic facts about points. It is natural, on the present approach, to reject the possibility of such facts, since the symmetry of GR tells us that they do no work in physical explanations. The principle of Leibniz equivalence thus extends straightforwardly to GR.

What's more, it permits the interpreter of GR to dodge an important problem the theory raises for naive forms of substantivalism. This is the "hole argument" (see Earman and Norton, 1987). If one does assume, contrary to the present approach, that there are distinct physical possibilities related by diffeomorphisms, a strange sort of unobservable indeterminism pops up in GR. This is because a diffeomorphism can sometimes leave a particular surface of simultaneity (a state of the world "at a time") unchanged while shuffling around what happens at which point in the future of that surface. Since the shuffling changes nothing invariant under GR's symmetries, the indeterminism disappears on the present approach, where only such invariants are physically real (Hofer, 1996).

Besides being unobservable, the hole-argument indeterminism is spooky in that it offers no probabilistic predictions about what will happen at which point. It simply entails that

many diffeomorphism-related solutions (which the believer in Leibniz equivalence would count as physically the same) are possible, without assigning likelihoods to any of them. If this sort of chanceless, unobservable indeterminism strikes us (as I think it should) as a theoretical problem for GR naively interpreted, the present approach to symmetries provides a solution.

Some foundational confusion arises, however, in applying the present approach to narrow down the fundamental physical quantities in GR. A natural requirement would seem to be that the fundamental quantities remain constant along “gauge orbits,” curves in state space that connect states related by diffeomorphisms. Quantities failing to meet this condition would appear to be altered by symmetry transformations. But as Earman (2002) argues, the list of quantities meeting this condition is very small, and includes no quantities that take on different values at different times. Does this mean there is never temporal change in any of GR’s fundamental quantities?

As Maudlin (2002) argues, this is an absurd result, since changing quantities like the position of Mercury are the source of our evidence for GR in the first place. So Earman’s argument must be an antinomy of sorts. Maudlin argues that this indicates a flaw in the present approach to symmetries, but an alternative (although closely related) diagnosis by Healey (2004) allows us to preserve the present approach while rejecting Earman’s conclusion. As Healey points out, Earman’s assumptions about what counts as invariant under the diffeomorphism symmetry of GR don’t hold up when the lesson of Leibniz equivalence is taken into account. A paradigm example of an observable quantity in GR is the scalar curvature of spacetime, $R(x)$. This quantity changes along a gauge orbit because the mathematical point x used to stand for a point p in physical spacetime can be altered by a diffeomorphism. So while R ’s value at p is left unchanged by the transformation, the formalism tells us (misleadingly) that since $R(x)$ is not unchanged, neither is the scalar curvature. This is the misstep in Earman’s antinomy.

The present approach to symmetry has acquitted itself admirably in the philosophy of space and time. It leads us to sensible and informative metaphysical conclusions while avoiding the troubling implications of Earman’s antinomy and the hole argument’s indeterminism. Moving on to other symmetries besides those of spacetime, we will find similar successes.

4 Global internal symmetries and Humeanism

The symmetry transformations we just discussed act on a system by taking some spatiotemporal distribution of physical structures and rearranging it into a different distribution of the same physical structures. There are also symmetries which involve only changes in the values of the physical quantities. Since these transformations alter the value of the quantities that constitute a physical structure, rather than changing anything spatially or temporally external to it, we call them “internal symmetries.”

The most familiar example arises in introductory electrostatics. The electrostatic potential $V(x)$ (measured in the unit of volts) is a quantity whose gradient at a point x determines the electric field, and therefore the force, at x . Since adding a constant to $V(x)$ (transforming $V(x) \rightarrow V(x) + c$) does not change its gradient, such a transformation makes no difference to the electrostatic force, and therefore no difference to the motion of charged particles. So any such transformation leaves the dynamical law (Coulomb’s law) unchanged, and is therefore a symmetry. We call it a “global” internal symmetry, since the transformation acts identically on the value of V at every point (i.e., globally).

In another important example, the wavefunction $\psi(x)$ which describes a system of particles in quantum mechanics is symmetric under a group of internal transformations called phase transformations. $\psi(x)$ is a complex-valued function on the $3N$ -dimensional space of possible configurations for N particles. The likelihood of the particles’ locations falling within a region R of this space on measurement is given by the integral of $|\psi|^2$ over this region. The square of a complex number is left unchanged if we multiply it by another complex number of absolute value one. The complex numbers of absolute value one, each of form $e^{i\theta}$, $\theta \in \mathbb{R}$ form a group called $U(1)$. Multiplying $\psi(x)$ at every point by the same element of $U(1)$ gives a phase transformation, which leaves all probabilities (and relative phases) unchanged and is therefore a symmetry of quantum mechanics.

Global internal symmetries of this sort present a challenge to a Humean combinatorial metaphysics where the fundamental quantities are point-sized.⁵ This challenge arises in cases where the symmetry transformations alter the numerical values of the fundamental point-sized quantities – which should lead us to expect that the representation of these quantities

⁵According to such a view (e.g. David Lewis’s “Humean supervenience”), all features of worlds like the one we live in supervene on the distribution of fundamental properties over points in spacetime.

contains some surplus structure. In some cases (the quantum wavefunction is one such) the surplus structure appears to include the representation of the quantities as localized at points. After “modding out” the surplus, we are left with quantities less local or point-sized than the ones we started with. Let’s look at the quantum example.

There are excellent reasons to be unsatisfied with a picture of quantum mechanics as describing only probabilities for particles to show up somewhere when measured. What counts as a measurement when all aspects of the world are presumably quantum? This is the aptly-named measurement problem. One possible route to a solution is to suppose that $\psi(x)$ is a physical field on configuration space (see Albert, 1996).

This step in solving the measurement problem runs directly afoul of the present approach to symmetries. The complex value of $\psi(x)$ is altered by phase transformations, which are symmetries of quantum theory. So therefore it cannot stand for a fundamental physical quantity, but must instead represent surplus structure.

$\psi(x)$ can’t be entirely surplus structure, since the values it takes at points fully determine a physical state and physical states must contain some qualitative information. So it must be that the wavefunction consists partly of real physical information, and partly of surplus structure. One way to get rid of the surplus is to suppose that $|\psi(x)|^2$ is physically fundamental, and the breakdown of ψ into real and imaginary components is surplus. Unfortunately, as Wallace (forthcoming, 52) points out, this leaves out information about the relative phases of ψ ’s components, which is of empirical import, since it determines the degree of interference between these components.

Alternative candidates for fundamental quantities are given by other phase-invariant quantities we can define in terms of $\psi(x)$. One such is the density operator $|\psi\rangle\langle\psi|$, an operator on Hilbert space which contains all of the information about the state included in the wavefunction. In a quantum theory defined on spacetime, like quantum field theory, one can also construct a “local” density operator describing the state’s behavior in any region of spacetime. Wallace and Timpson (forthcoming) have advanced an attractive view they call “spacetime state realism,” according to which the density operators assigned to regions stand for the fundamental quantities.

Due to the existence of entangled states, which exhibit instantaneous correlations between distant parts of a physical system, the density operators which are fundamental in spacetime state realism do not conform to Humean combinatorial metaphysics. The density operator

of the region made up of regions A and B (their union) is not uniquely determined by the operators assigned to A and B . In this way, a “spacetime state” universe is not built up of localized parts that fully determine its large-scale features. Spacetime state realism, favored by the present approach over the alternative wavefunction realist view, differs from that view in part by denying what Lewis called Humean supervenience.

There may be alternative Humean ontologies for quantum theory which do respect phase symmetry. For instance, Tumulka (2009) has proposed an ontology of instantaneous point-sized “flashes” which are obviously invariant. But Tumulka’s ontology has so far been applied to only a few toy quantum theories, whereas spacetime state realism applies to all of them (since it uses only the pre-existing quantum formalism).

5 Gauge potentials and Aharonov-Bohm

When I transform the phase of a wavefunction $\psi(x)$, I do so by multiplying it by the same unit complex number at all points x . It is also possible for a similar theory to be symmetric under internal transformations that differ from point to point (and even time to time). This more complex type of symmetry is called a “local” internal symmetry, or local gauge symmetry. Theories exhibiting it are called “gauge theories.”

Such theories make additional trouble for the Humean, if we hold to the present approach. They also provide a case (like that of general relativity) in which the present approach dodges serious problems of interpretation having to do with determinism. For in gauge theories, ontologies which eschew the present approach to preserve a metaphysics of Humean point-sized quantities face the problem of chanceless, unobservable indeterminism.

Although they may sound exotic, gauge theories are everywhere. The best-known example is classical electrodynamics. In relativistic electrodynamics, the vector potential $A_i(x)$ is an often-used quantity. As a relativistic invariant, it is a four-dimensional vector; the time component is the familiar electric potential $V(x)$ (discussed above) while the space components collectively form the magnetic potential. Since the value of $V(x)$ is physically unimportant, we would expect the same to be true of $A_i(x)$ ’s components, and indeed that is correct. Physical predictions are fixed by the field tensor $F_{ij}(x) = \partial_i A_j(x) - \partial_j A_i(x)$. This quantity is left unchanged if we add the gradient of a scalar field to $A_i(x)$, transforming $A_i(x) \rightarrow A_i(x) + \partial_i \Lambda(x)$. These are called (local) gauge transformations – local since they

take on different values at different points x . Adding a gauge transformation to a vector potential solution of Maxwell's equations always gives us another solution.

Nothing prevents $\Lambda(x)$ from taking on non-zero values in some parts of spacetime but not others. In particular, in a given frame $\Lambda(x)$ might be zero prior to t and non-zero after t . This means that two solutions to Maxwell's equations – two possibilities, according to electrodynamics – can agree completely about the potential up until t while disagreeing about its values thereafter. On any reasonable definition of determinism, this means that $A_i(x)$ does not evolve deterministically – its evolution is not even predicted probabilistically. It is therefore a very poor choice of fundamental quantity for electromagnetism, which is normally taken to be deterministic and whose observables evolve deterministically (except in certain strange cases).

Cast in the extremely general “fiber bundle” formalism, this indeterminism is analogous to the “hole problem” in general relativity we discussed earlier (see Healey, 2007). The solution ought to be similar, i.e. a choice of potential should be similar to a choice of coordinates, a conventional way of speaking about fundamental reality. One easy way to get this result is just to stipulate that the field tensor $F_{ij}(x)$ is fundamental. (Since this quantity isn't changed by gauge transformations, we call it “gauge invariant.”) But such a posit is undermined by a phenomenon called the Aharonov-Bohm effect.

In an Aharonov-Bohm experiment, a current is sent through a long (represented as infinite) solenoid, or coiled electromagnet. The field $F_{ij}(x)$ remains zero outside the solenoid, but the phase of complex waves passing by the solenoid is nonetheless altered. Since the wavefunction $\psi(x)$ of a quantum particle is a complex wave, and since its differences in phase from other particles' wavefunctions is observable, this leads to a measurable difference in the behavior of quantum particles passing by solenoids despite the zero field in their vicinity. If electromagnetic fields are the explanation for this behavior, they must act on quantum particles at a distance – not a very elegant picture.

There are other gauge-invariant ontologies that can explain Aharonov-Bohm without action at a distance, but as with spacetime state realism they are incompatible with a Humean metaphysics. On one view, the fundamental quantities are so-called “holonomies,” which are defined not at points but instead on closed loops in spacetime (Healey, 2007). Obviously these are not local quantities in the sense preferred by Humean supervenience. On another view, the fundamental quantities are gauge-invariant values of the so-called “connection,”

which are (in Maudlin’s parlance) “hyper-local,” so that there is no determinate matter of fact about whether distant spacetime points agree as to their value (Maudlin, 2007, 78-103). Either way the present approach to symmetries, combined with the reasonable physicist’s bias against discontinuous action at a distance, dictates a surprisingly revisionist ontology – but one that preserves determinism against the specter of spooky unobservable indeterminism.

6 Further questions

Considerations of space prevent me from posing, let alone addressing, all the interesting problems in this area. To what extent do symmetries dictate the nature of forces? (Martin, 2002) How can they have empirical import? (Healey, 2009) Is the group a general enough concept to represent all symmetries? (Guay and Hepburn, 2009) I leave the pursuit of these problems to the interested reader.

Acknowledgements

Thanks to David Manley, Sarah Moss and Eric Swanson for valuable comments and discussion on a previous draft and to Gordon Belot for numerous illuminating comments and discussions.

References

- Albert, David Z. (1996), “Elementary Quantum Metaphysics,” in Cushing, James T., Arthur Fine and Sheldon Goldstein (eds.), *Bohmian Mechanics and Quantum Theory: An Appraisal*, Netherlands: Kluwer, 277–284.
- Belot, Gordon (2001), “The Principle of Sufficient Reason,” *Journal of Philosophy* 98:55–74.
- Belot, Gordon (in progress), “A Didactic Rant, or Of Symmetry,” .
- Earman, John (2002), “Thoroughly Modern McTaggart,” *Philosophers’ Imprint* 2:<http://www.philosophersimprint.org/002003/>.
- Earman, John and John Norton (1987), “What Price Spacetime Substantivalism? The Hole Story,” *British Journal for the Philosophy of Science* 38:515–525.

- Guay, Alexandre and Brian Hepburn (2009), “Symmetry and Its Formalisms: Mathematical Aspects,” *Philosophy of Science* 76:160–178.
- Healey, Richard (2004), “Change Without Change, and How to Observe it in General Relativity,” *Synthese* 141:381–415.
- Healey, Richard (2006), “Symmetry and the Scope of Scientific Realism,” in Demopoulos, William and Itamar Pitowsky (eds.), *Physical Theory and its Interpretation*, Netherlands: Springer, 143–160.
- Healey, Richard (2007), *Gauging What’s Real*, Oxford: Oxford UP.
- Healey, Richard (2009), “Perfect Symmetries,” *British Journal for the Philosophy of Science* 60:697–720.
- Hofer, Carl (1996), “The Metaphysics of Spacetime Substantivalism,” *Journal of Philosophy* 93:5–27.
- Ismael, Jenann and Bas C. van Fraassen (2003), “Symmetry as a Guide to Superfluous Theoretical Structure,” in Brading, Katherine and Elena Castellani (eds.), *Symmetries in Physics: Philosophical Reflections*, Cambridge: Cambridge UP, 371–392.
- Ladyman, James and Don Ross (2007), *Every Thing Must Go*, Oxford: Oxford UP.
- Lewis, David (1970), “How to define theoretical terms,” *Journal of Philosophy* 67:427–446.
- Martin, Christopher (2002), “Gauge Principles, Gauge Arguments and the Logic of Nature,” *Philosophy of Science (Proceedings)* 69:S221–S234.
- Maudlin, Tim (2002), “Thoroughly Muddled McTaggart,” *Philosophers’ Imprint* 2:<http://www.philosophersimprint.org/002004/>.
- Maudlin, Tim (2007), *The Metaphysics Within Physics*, Oxford: Oxford UP.
- Roberts, John T. (2008), “A Puzzle about Laws, Symmetries and Measurability,” *British Journal for the Philosophy of Science* 59:143–168.
- Tumulka, Roderich (2009), “The Point Processes of the GRW Theory of Wave Function Collapse,” *Reviews in Mathematical Physics* 21:155–227.
- Wallace, David (forthcoming), “Quantum Measurement Problem: State of Play,” in Rickles, Dean (ed.), *Ashgate Companion to the New Philosophy of Physics*, Ashgate.
- Wallace, David and Christopher Gordon Timpson (forthcoming), “Quantum Mechanics on Spacetime I: Spacetime State Realism,” *British Journal for the Philosophy of Science* .