
SYMMETRY AND EQUIVALENCE

Gordon Belot * †

1 Introduction

My topic is the relation between two notions, that of a symmetry of a physical theory and that of the physical equivalence of two solutions or models of such a theory. In various guises, this topic has been widely addressed by philosophers in recent years.¹

As I intend to use the term here, a symmetry of a theory is a map or transformation that leaves invariant the structure used to encode the laws of the theory. The notion of physical equivalence that I have in mind is as follows: two solutions (models) of a physical theory are *physically equivalent* if and only if, for each possible physical situation, the two are equally well- or ill-suited to represent that situation. The first of these notions is a formal one, in the sense that in specifying the formalism of a theory, one specifies its symmetries. The second is an interpretative notion: two solutions that are

*To appear in R. Batterman (ed.), *The Oxford Handbook of Philosophy of Physics*.

†Versions of this paper were given (under a different and better title) in Chicago, Ann Arbor, and Montreal—thanks to all those present. For helpful comments and discussion, thanks to Dave Baker, Bob Batterman, Katherine Brading, Caro Brighthouse, Kevin Coffey, Hilary Greaves, Juha Pohjanpelto, Peter Railton, Laura Ruetsche, Jamie Tappenden, and Bas van Fraassen. Special thanks to Old Man Batterman for his great patience and editorial sagacity—even if he doesn't know a good title when he sees one.

¹See, e.g., Baker (2010), Brading and Castellani (2007), Dasgupta (2010), Debs and Redhead (2007), Healey (2009), Ismael and van Fraassen (2003), North (2010), and Roberts (2008).

physically equivalent according to one interpretation of a given formalism may be inequivalent according to another.

Part of the interest in these notions among philosophers derives from the following line of thought. In an influential discussion John Earman argued that in the special case of spacetime symmetries, facts about symmetries place interesting constraints on good interpretative practice.² Consider the case of time translation by b temporal units. In the first instance, this operation is a map that takes each point of spacetime to a point of spacetime b units later than itself. But we can also think of time translation as acting in an obvious way on fields living on spacetime. Earman observes that in this context, a transformation like time translation is a symmetry of a theory if and only if it maps solutions of the theory to solutions of the theory. And he goes on to argue that under any reasonable interpretation, a transformation of this kind ought to be a symmetry of the theory if and only if it is a symmetry of the geometric structure of spacetime.³ In particular, it follows from Earman's arguments that if one focuses on theories in which spacetime symmetries are the only relevant symmetries, then under any reasonable interpretation, solutions are physically equivalent if and only if they are related by a symmetry—in this way formal facts place interesting constraints on (good) interpretation.

It is natural to wonder how much of this picture carries over if one moves beyond spacetime symmetries. Much recent philosophical writing on symmetries and related topics seems to assume that the answer is: all of it.⁴ For the following two doctrines play an important (if often implicit) role in this literature.

D1. The symmetries of a classical theory are those transformations that

²See Earman (1989, §3.4).

³If spacetime geometry fails to be invariant under some such transformation that is symmetry of the laws, then one is positing unnecessary geometric structure (think of Newton on absolute space). Conversely, if one's spacetime structure is invariant under transformations that are not symmetries of the laws, then one is cheating somehow or other, employing some structure in one's formalism to break a spacetime symmetry without being willing to pay the ontological price (think of someone whose theory has a point mass springing back and forth about the origin in the Cartesian plane, but who is unwilling to posit a privileged point of space or any physical structure—matter, field, force, etc.—other than the single moving point-particle).

⁴I won't do anything to substantiate this claim here (let us not descend into recriminations at this time!), beyond saying that, like many a rant, this one is directed in part at earlier versions of its author.

map solutions of the theory's equation of motion to solutions of the theory's equation of motion.

- D2.** Two solutions of a classical theory's equation of motion are related by a symmetry if and only if they are physically equivalent, in the sense that they are equally well- or ill-suited to represent any particular physical situation.

Here D1 is a formal condition, D2 a constraint on good interpretative practice. Combined, they would yield a strong and interesting constraint on (good) interpretation.

However, it is not difficult to see that the combination of D1 and D2 is an unhappy one. Consider any classical theory. Let u_1 and u_2 be solutions of the theory's equation of motion. Consider the transformation T on the space of solutions that maps u_1 to u_2 , u_2 to u_1 , and leaves every other solution where it is. According to the first doctrine above, T is a symmetry of the theory. The second doctrine above then implies that u_1 and u_2 are physically equivalent. So the two doctrines jointly imply that any two solutions of any classical theory are physically equivalent—and hence that each classical theory is more or less useless because unable to discriminate among the systems for which it provides good models.⁵

There are many strategies for encoding the laws of a theory in a mathematical form—and each such strategy is associated with a formal notion of symmetry. D1 encapsulates the notion of symmetry associated with the most spare and naïve method of encoding the laws of a theory. It is natural to wonder whether some more sophisticated relative of D1 might form a fruitful combination with D2. There has been some recent discussion of this question among philosophers, with some authors adopting an optimistic view, arguing or assuming that some more sophisticated relative to D1 will do the trick, and some authors adopting a pessimistic view, according to which the notion of physical equivalence is not closely linked with a formal notion of symmetry (even when we restrict to well-behaved interpretations).⁶

⁵Our world contains many systems that can be treated (for all practical purposes) as being isolated. It follows from the conjunction of D1 and D2 that if a theory provides equally good models of two isolated systems, then *any* model of the theory must provide equally good representations of each of these systems—which is to say that the theory is just about useless in application.

⁶Since none of the authors I have in mind address the topic in quite the same terms

My project here is to examine the viability of D1 and D2 and to say a bit about the source of their *prima facie* plausibility. In the next section, I briefly sketch the framework in which I will proceed. Section 3 is devoted to a discussion of D1 and its inadequacies. Sections 4 and 5 give a brief overview of some of the more sophisticated alternatives to D2. Section 6 argues that none of the notions of symmetry considered combine satisfactorily with D2. Section 7 presents my cautiously pessimistic conclusions.

2 Stage-Setting

Consider a classical theory in which a history of a physical system is represented by a function $u : M \rightarrow W$; here M is the manifold parameterised by the independent variables of the theory and W is the manifold parameterised by the dependent variables of the theory.⁷ Thus M might correspond to time and W to the space of possible configurations of some system of particles. Or M might correspond to spacetime and W to the space of values available to some field at each point of spacetime. The laws of the theory are given by an equation,

$$\Delta[u] = 0,$$

employed here (or in quite the same terms as one another), the following attributions should perhaps be taken with a grain of salt. On the optimistic side: Baker (2010, §1) conjectures that D2 is true relative to some generalization of the Hamiltonian approach discussed in §5 below; Brading and Castellani (2007, §§4.1 and 8.2) claim that D2 is true relative to the notion of a classical symmetry discussed in §4 below; Roberts (2008, fn. 3) holds that D2 is true for the sort of version of D1 discussed in fn. 11 below. On the pessimistic side: Ismael and van Fraassen (2002) take D1 to be the only philosophically salient notion of symmetry—and thus think that in order to get around the sort of problems noted above, one should take models to be physically equivalent if and only if they are related by a symmetry and agree in (roughly speaking) their directly perceivable features; Healey (2009) likewise notes some of the problems surrounding D1, mentions that there are other formal notions of symmetry available, then pursues an approach that builds physical equivalence into its notion of symmetry; Dasgupta (2010) argues that a notion of symmetry appropriate to D2 must involve notions like that of a mental state.

⁷More generally, here and throughout, one could take histories to be represented by sections of a fibre bundle $E \rightarrow M$ with typical fibre W . Almost any classical field can be thought of this way. E.g., a gauge field, standardly represented by a connection one-form on a principle bundle $P \rightarrow M$, can be represented instead by a section of a certain affine bundle $J^1(P)/G \rightarrow M$ associated with P ; see, e.g., Kolář *et al.* (1993, §17.4).

where Δ is a (linear or nonlinear) differential operator.⁸ The job performed by such an equation is to single out from a large space \mathcal{K} of functions that represent “kinematically possible” histories of the system those that represent dynamically possible histories. Typically, \mathcal{K} is specified by specifying the independent and dependent variables of the theory and the degree of regularity (smoothness etc.) of candidate functions as well as any (asymptotic) boundary conditions that they must satisfy. We call the subspace $\mathcal{S} \subset \mathcal{K}$ consisting of u that satisfy the equation of the theory the *space of solutions*.

3 A Recipe for Disaster

Here is a recipe for arriving at the characterization of symmetries of classical theories that is embodied in the doctrine D1 discussed above: begin with the standard abstract notion of a symmetry; reflect on the case of spacetime symmetries; extrapolate.

SYMMETRIES, ABSTRACTLY SPEAKING. There is a basic notion of symmetry, elaborated in various ways in the various branches of mathematics. Consider a *structure*—a set of objects, D , equipped with some relations, R_1, \dots , and functions, f_1, \dots . A *symmetry* of this structure is a one-to-one and onto map $F : D \rightarrow D$ that preserves all of the relations and functions between the objects.⁹

SPACETIME SYMMETRIES. Suppose that we want to check whether a given field theory is invariant under a spacetime transformation such as time-translation. Then all we have to do is to check to make sure that whenever $u(x, t)$ is a solution of the theory’s equation of motion, so is $u(x, t - b)$, for any real number b . That is, we define in the obvious way an operator on the set of kinematically possible fields that implements the putative symmetry, then check to see whether it maps solutions to solutions.

⁸So the left-hand side of this equation is the function on M that results from applying Δ to u and the right-hand side is the zero function.

⁹I.e., $R(x_1, \dots, x_n)$ if and only if $R(F(x_1), \dots, F(x_n))$; and $f(x_1, \dots, x_n) = y$ if and only if $f(F(x_1), \dots, F(x_n)) = F(y)$.

GENERALIZATION. Suppose that we think of a classical theory as a structure consisting of a large set of functions \mathcal{K} with a distinguished subset \mathcal{S} . Then we would expect the symmetries of such a theory to be those one-to-one and onto maps from \mathcal{K} to itself that preserve \mathcal{S} . Spacetime symmetries fit this pattern exactly! This suggests that we ought to think of classical theories in the way just outlined.

SOMETHING HAS GONE WRONG. Following this recipe leads to the following

Fruitless Definition: A *symmetry* of a differential equation Δ is a one-to-one and onto map $T : u \in \mathcal{K} \mapsto T(u) \in \mathcal{K}$ that preserves the space of solutions, in the sense that $u \in \mathcal{S}$ if and only if $T(u) \in \mathcal{S}$.

As an attempt to capture the ordinary notion of a symmetry of an equation (and hence of a theory), this is a disaster. Ordinarily, symmetries of theories are hard to come by. But some remarkable theories have atypically large symmetry groups (relative to the number of degrees of freedom of the systems that they treat). The definition above effaces this sort of distinction between theories. For if we allow arbitrary permutations of the solutions of a theory to count as symmetries, then the size of a theory's group of symmetries depends only on the size of its space of solutions.¹⁰

All of this is liable to strike fans of the Fruitless Definition as cheap and misleading: after all, since the spaces involved have topological and differential structure, surely one should restrict attention to continuous or smooth transformations of the space of kinematic possibilities? This is fair enough—but it wouldn't make much difference. At best, it would lead us

¹⁰Don't mathematicians sometimes offer up characterizations of symmetries along these lines? Yes—but only when speaking loosely and heuristically. Thus in the introduction to Olver's influential textbook on symmetries of differential equations, we are told that: "Roughly speaking, a symmetry group of a system of differential equations is a group which transforms solutions of the system to other solutions;" Olver (1993, xviii)—see also, e.g., Bluman and Anco (2002, 2) and Klainerman (2008, 457). But on the next page we are told that once "one has determined the symmetry group of a system of differential equations, a number of applications become available. To start with, one can directly use the defining property of such a group and construct new solutions to the system from known ones." Of course, one cannot do this if one's notion of a symmetry is given by the Fruitless Definition—one needs to be working with one of the more specialized notions that are the focus of Olver's book, some of which are described below. And likewise for the other applications on Olver's list.

to identify the symmetries of a theory with something like the family of smooth permutations of the space of solutions, rather than with the full family of permutations of that space.¹¹ But this would still lead us to the conclusion that every theory had an enormous group of symmetries that depended only on relatively coarse features of its space of solutions.¹²

4 Symmetries of Differential Equations

Laws of classical physics are given by differential equations—and there are a number of ways that a differential equation can be encoded in a structure whose symmetries can be investigated. D1 corresponds to the most flat-footed of these. In this section and the next we briefly survey some more sophisticated alternatives: the focus in this section is on approaches that directly encode the equation of motion of the theory in a structure whose symmetries can be identified with the symmetries of the theory; in the next section, we consider approaches that take a detour via a Lagrangian or Hamiltonian formulation of the theory.

There is no settled, definitive notion of a symmetry of a differential equation—rather there are a family of related notions.¹³ Three are espe-

¹¹ Roughly speaking: for non-freaky Δ , one expects \mathcal{S} to sit inside \mathcal{K} as (something like) a submanifold; if ϕ_1 and ϕ_2 are (generic) points in \mathcal{S} then there will be a diffeomorphism $F : \mathcal{S} \rightarrow \mathcal{S}$ with $\phi_2 = F(\phi_1)$; and one expects to be able to extend F to a diffeomorphism $\bar{F} : \mathcal{K} \rightarrow \mathcal{K}$. (Slightly more carefully: one expects that any obstructions to extending F in this way are going to be technical in nature, and not detract from the main point.)

¹²Consider, by way of illustration, the Newtonian theory of three gravitating point particles of distinct masses. Here a point in the space of kinematically possible fields, \mathcal{K} , essentially assigns each of the particles a worldline in spacetime (without worrying about whether these worldlines jointly satisfy the Newtonian laws of motion). The space of solutions, \mathcal{S} , is the eighteen-dimensional submanifold of \mathcal{K} consisting of points corresponding to particle motions obeying Newton's laws. So one expects that any diffeomorphism from \mathcal{S} itself can be extended to a suitably nice map from \mathcal{K} to itself. But for any solutions $u_1, u_2 \in \mathcal{S}$, we can find a diffeomorphism from \mathcal{S} to itself that maps u_1 to u_2 , so we again find that arbitrary pairs of solutions are related by symmetries. This seems unacceptable, since we ordinarily think of this theory as having a relatively small symmetry group (consisting just of spacetime symmetries).

¹³For an historical overview and references, see Olver (1993, 172 ff. and 374 ff.). There is of course a trade-off to be made between fecundity and generality. In practice, the notions that mathematicians and physicists find interesting are far more restrictive than the fruitless notion considered above. But there would appear to be no feeling that

cially important: the notion of a classical symmetry, the notion of a generalized symmetry, and the notion of a non-local symmetry. I will sketch the first of these and then describe how the other two are related to it. A few more details concerning classical symmetries are presented in an appendix below.

Let us focus on the field-theoretic case: the independent variables of the theory parameterise the spacetime manifold M ; a configuration of the field is represented by a function $u : M \rightarrow W$; and the dynamical laws of the theory are encoded in a k^{th} -order partial differential equation Δ .

Roughly speaking, the classical symmetries of such a theory can be characterized as follows. Let E be the manifold $M \times W$ that is parameterised by the independent and dependent variables of the theory taken together. Any diffeomorphism $\bar{d} : E \rightarrow E$ will induce a transformation d from the space of kinematically possible fields of the theory to itself (this claim is unpacked in the appendix below). A map d that arises in this way is a *classical symmetry* of the equation if and only if it maps solutions to solutions.¹⁴ So classical symmetries are, roughly speaking, transformations of the space of kinematically possible fields that: (a) map solutions to solutions; and (b) are suitably local, in the sense that they arise from smooth transformations of the dependent and independent variables of the theory. Classical symmetries are also known as *Lie symmetries* or *point symmetries*. The *spacetime symmetries* of a theory are those classical symmetries that arise from transformations of the independent variables of the theory that leave the dependent variables untouched.¹⁵

The classical symmetries of an equation form a family *much* more restricted than that picked out by D1—and typically this means that solutions related by a classical symmetry share many salient features. In typical cases, the classical symmetries of a theory of n Newtonian particles are just the spacetime symmetries—so solutions are related by a symmetry if and only if

there is a *correct* notion of a symmetry of a differential equation—plausibly, for every interesting such notion, there is a yet more general one that is still interesting.

¹⁴Two subtleties are glossed over in the text. (i) In general d may map a kinematically possible field to an object that is a sort of partially-defined multiply-valued field rather than a kinematically possible field (see fn. 60 below). (ii) The story is (yet) more complicated in the important special case where $\dim W = 1$.

¹⁵Note, in particular, that we make no appeal here to metric tensors and the like—we just look for transformations of the independent variables that leave invariant the equations of motion. (Some subtleties arise at this point for certain types of fields; see, e.g., Kolář *et al.* (1993) on gauge natural bundles.)

they have the particles instantiating the same sequence of relative distance relations.¹⁶ In the case of general relativity, the classical symmetries are generated by spacetime diffeomorphisms and by scale transformations—so that, roughly speaking, two solutions are related by a symmetry if and only if they agree about the pattern of ratios of distances instantiated.¹⁷

The notion of a classical symmetry is a special case of the basic mathematical notion of symmetry: the classical symmetries of a differential equation are the structure-preserving maps for a certain structure associated with the equation. Here is the story in briefest outline (a few more details are provided in the appendix below).¹⁸ One constructs a space, $J^k(E)$ (the k^{th} jet bundle over E) a point of which is specified, intuitively speaking, by specifying a point x of spacetime and the values of the field and its partial derivatives through order k at x . $J^k(E)$ comes equipped with some geometric structure, \mathcal{C} (the *Cartan distribution*), that ensures that the specification of a point of $J^k(E)$ lives up to this intuitive picture. Our differential equation Δ determines a submanifold $\mathcal{E} \subset J^k(E)$ —since for each point of spacetime Δ imposes a constraint on the values of the fields and their partial derivatives at that point. So we have a structure consisting of a large space, $J^k(E)$, a geometric structure \mathcal{C} on $J^k(E)$, and a distinguished subspace \mathcal{E} of $J^k(E)$. In accord with our template for defining symmetries of structures, we take a symmetry of this contraption to be a one-to-one and onto map $F : J^k(E) \rightarrow J^k(E)$ that respects \mathcal{C} and maps points of \mathcal{E} to points of \mathcal{E} . It turns out that the resulting notion coincides with the notion of a classical symmetry as characterized above.

In order to characterize generalized and non-local symmetries, it is helpful (as well as more honest—see fn. 60 below) to describe the relevant notions by characterizing their infinitesimal generators.¹⁹

From this perspective, classical symmetries (investigated by Sophus Lie and others in the late nineteenth century) are, roughly speaking, transformations that map solutions to solutions and whose infinitesimal generators depend only on the independent and dependent variables of the theory.

Emmy Noether introduced a substantially more general notion at the

¹⁶ Hydon (2000, example 4.4).

¹⁷See Anderson and Torre (1996).

¹⁸For what follows, see, e.g., Krasil'shchik and Vinogradov (1999, Chapter 3).

¹⁹Unfortunately, discrete symmetries threaten to go missing when one adopts this perspective. For one approach to this problem, see Hydon (2000, Chapter 11).

beginning of the twentieth century.²⁰ Her *generalized symmetries* (also known as *local* or *higher* symmetries or as *Lie–Bäcklund* transformations) are, roughly speaking, transformations that map solutions to solutions and whose infinitesimal generators depend on derivatives of the fields as well on the independent and dependent variables of the theory.²¹ These arise as the symmetries of a fancier structure that can be used to encode the differential equation of interest: again the equation is represented as a submanifold of an ambient space equipped with a Cartan structure; this time the ambient space is an *infinite* jet bundle—the specification of a point of which involves specifying a point of spacetime together with the values at that point of a function and *all* of its partial derivatives.

Every classical symmetry is also a generalized symmetry. Some equations have no non-classical generalized symmetries (the Einstein field equations of general relativity are an example).²² But many equations have generalized symmetries that are not classical symmetries. Striking examples include:

The Kepler Problem. This is the problem of determining the motion of Newtonian massive particle moving in the external gravitational field of a fixed massive body.²³ In addition to the ‘obvious’ conserved quantities for this system—energy and angular momentum—there is a ‘hidden’ one, the Lenz–Runge vector.²⁴ Associated with this hid-

²⁰See, e.g., Olver (1993, Chapter 5) or Krasil’shchik and Vinogradov (1999, Chapter 4).

²¹To get a feeling for what this means, consider the sort of gauge transformations that normally arise in presentations of Maxwell’s theory: if we take the vector potential $A(x)$ as our field, then the theory is invariant under infinitesimal transformations of the form $A \mapsto A + \varepsilon d\Lambda$ where $\Lambda(x)$ is a real-valued function on spacetime. Now suppose that Λ is a map then when fed a kinematically possible A returns a real-valued function $\Lambda[A]$ on spacetime. If for each spacetime point x and each A , the value of $\Lambda[A]$ at x depends only on x and on $A(x)$, then the infinitesimal transformation $A \mapsto A + \varepsilon d\Lambda[A]$ corresponds to a classical symmetry of Maxwell’s theory; if the value of $\Lambda[A](x)$ depends also on a finite number of derivatives of A at x , then this map is a generalized symmetry of Maxwell’s theory. See the discussion of generalized gauge symmetries in Pohjanpelto (1995) and in Torre (1995). For a thoroughly worked-out example involving only finitely many degrees of freedom, see Cantwell (2002, §14.4.1).

²²See Anderson and Torre (1996).

²³Note that the Kepler problem contains all of the dynamics of the honest Newtonian two-body problem. See, e.g., Goldstein *et al.* (2002, §3.1).

²⁴The Lenz–Runge vector is $-\frac{q}{|q|} + \frac{1}{\mu}p \times (q \times p)$, where q is the position of the moving particle, p is its momentum, and μ is a constant that depends on the masses.

den conserved quantity is a non-classical generalized symmetry of the Kepler problem (more on this below).²⁵

The KdV Equation. The *Korteweg–de Vries equation* is an equation used to model waves in shallow water. We can think of it as an equation governing a field $u(x, t)$ on a two-dimensional spacetime:

$$\frac{\partial u}{\partial t} + \frac{\partial^3 u}{\partial x^3} + u \frac{\partial u}{\partial x} = 0.$$

The only classical symmetries of this equation are spacetime and scaling symmetries; but it has an infinite-dimensional family of generalized symmetries.²⁶

Yet more general are the non-local symmetries that have been investigated in recent decades.²⁷ These are, roughly speaking, transformations that map solution to solutions and whose infinitesimal generators are allowed to depend on non-local functionals (such as integrals) of the fields, as well as on their derivatives and on the dependent and independent variables of the equation.²⁸ They arise as symmetries of a structure that results when the differential equation is encoded in a certain extension of the infinite jet bundle. Many examples of non-local symmetries are known—even the Kepler problem admits nonlocal symmetries that are not generalized symmetries.²⁹

5 Variational and Hamiltonian Symmetries

The governing equations of many physical theories can be given a Lagrangian or Hamiltonian treatment (although this sometimes requires some

²⁵For the generalized symmetry of the Kepler problem, see Lévy–Leblond (1971, §5.B). Note that this system also admits a classical symmetry that is in a sense associated with the Lenz–Runge vector; see Prince and Eliezer (1981).

²⁶For the classical and generalized symmetries of the Korteweg–de Vries equation, see Olver (1993, 125 ff. and 312 ff.).

²⁷See, e.g., Krasil’shchik and Vinogradov (1999, Chapter 6).

²⁸For a thoroughly worked-out example in which an infinitesimal symmetry and the corresponding finite symmetry both depend on an integral over space of the dependent variable of the theory, see Cantwell (2002, §16.2.2.1).

²⁹See, e.g., Leach *et al.* (2003).

craft or flexibility). For such theories, it is natural to consider the symmetries of the structures employed in the Lagrangian or Hamiltonian treatment—which need not coincide with the symmetries of the equation itself.

VARIATIONAL SYMMETRIES. There there are two main styles of Lagrangian formalism, sometimes called the *dynamical* and the *covariant* approaches, which coincide for theories with finitely many degrees of freedom but differ for field theories.³⁰ Dynamical approaches are quite similar in spirit to the Hamiltonian approach considered below.³¹ Under covariant approaches, which are our concern here, the basic space of interest is a jet bundle over the space of dependent and independent variables of the theory (as in the treatment of classical symmetries sketched above) and the Lagrangian is now to be thought of as an object that when fed a kinematically possible field u gives in return a d -form $L[u]$ on the spacetime M (here $d = \dim M$).³² Each such Lagrangian is associated with a differential equation, whose solutions correspond to those fields that satisfy a certain variational principle for the Lagrangian. A (classical or generalized) symmetry of an equation arising from a Lagrangian is called a *variational symmetry* if it leaves the Lagrangian invariant. Even if a theory admits a Lagrangian formulation, some symmetries of the underlying equation may not show up as variational symmetries—as a rule, Galilean boosts and scaling symmetries are often not variational.³³ Modulo niceties about boundary conditions, Noether’s theo-

³⁰For these labels and for an investigation of the relation between the two approaches, see Castrillón López and Marsden (2008). For introductions to the two approaches, see Abraham and Marsden (1985, §§3.5 ff. and Example 5.5.9).

³¹Under such approaches, one works with a space of instantaneous states (which will be infinite-dimensional in the field-theoretic case), equips this space with a real-valued function, L (the Lagrangian), and employs a variational principle to find those curves in the space of states that correspond to dynamically possible histories of the system.

³²For instances of this styles of approach, see, e.g., Castrillón López and Marsden (2008) and Krasil’shchik and Vinogradov (1999, Chapter 5).

³³On scale transformations, see Olver (1993, 255). A standard remedy is to introduce the notion of a *divergence symmetry*, a transformation that leaves the Lagrangian invariant up to a total divergence; many interesting symmetries are divergence symmetries but not variational symmetries, including boosts of Newtonian systems and the conformal symmetries of the wave equation; see Olver (1993, 278–281). Scaling symmetries are more subtle. Scale transformations are symmetries of general relativity, but are neither variational nor divergence symmetries; see Anderson and Torre (1996, §2.B). Rescaling of space and time is a symmetry of the wave equation that is neither a variational nor a divergence symmetry, although there is a related scale transformation that acts on the

rem assures us that each one-parameter family of variational symmetries is associated in a canonical way with a conservation law.³⁴

HAMILTONIAN SYMMETRIES. Let us restrict attention to the ideal case in which specifying an initial data set (=a possible instantaneous dynamical state) for the equation of the theory determines a unique solution defined at all times. Then, speaking roughly and heuristically, a Hamiltonian treatment amounts to the following. The *phase space* of the theory is the space, \mathcal{I} , of all initial data sets. The *Hamiltonian*, H , of the theory is the real-valued function on the phase space that assigns to each point of the phase space the energy of the corresponding physical state. The phase space can be equipped with a geometric structure, ω , with the following marvellous feature: together H and ω determine a family of curves in \mathcal{I} , exactly one passing through each point; each of these curves corresponds to a solution of the theory's equation of motion, in the sense that that the two objects pick out the same sequence of instantaneous dynamical states; and for any such solution there is a corresponding curve of this kind.³⁵ So the structure (\mathcal{I}, ω, H) in effect encodes the differential equation of the theory. It is natural to investigate the *Hamiltonian symmetries* of the theory: those one-to-one and onto maps from \mathcal{I} to itself that preserve both ω and H .³⁶

dependent variables as well as the independent variables which is a divergence symmetry (but not a variational symmetry); see Olver (1993, Examples 2.43, 4.15 and 4.36).

³⁴But: certain types of variational (or divergence) symmetries of theories whose initial value problems are ill-posed are associated with so-called trivial conservation laws; see Olver (1993, 342–346) on Noether's second theorem. And: there exist techniques for associating conservation laws with symmetries that do not rely on Noether's theorem; see, e.g., Bluman (2005).

³⁵ ω is a *symplectic form*—a closed, nondegenerate two-form. ω and H determine a vector field X_H on \mathcal{I} : X_H is the vector field that when contracted with ω yields the one-form dH . Integrating this vector field gives the curves mentioned in the text. Note that there is a canonical recipe for constructing \mathcal{I} , H , and ω given a Lagrangian treatment of the theory.

³⁶Note that the symplectic space (\mathcal{I}, ω) has a vast family of symmetries. Suppose that we are interested in a Newtonian theory of finitely many particles. Then \mathcal{I} is finite-dimensional, but the family of smooth permutations of \mathcal{I} that preserve ω is infinite-dimensional—it is only when we restrict attention to transformations that also preserve H that we end up with something like what we want. Something similar is of course true in ordinary quantum mechanics: while the family of unitary transformations of a Hilbert space will be very large, the family of such transformations that preserve a given Hamiltonian will be quite small—and only the latter is a good candidate for the

The Hamiltonian version of Noether’s theorem assures us (fussing about boundary conditions aside) that there is a conserved quantity associated with each one-parameter family of Hamiltonian symmetries of a theory.

6 Symmetry and Physical Equivalence

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—boosts are symmetries of his laws of motion but not of the spatiotemporal structure that he postulated. Reflection on this example and others motivates the principle that sound interpretative practice requires that the spacetime symmetries of one’s ontology ought to include the spacetime symmetries of one’s preferred theory. Conversely, it seems like something has gone wrong if the spacetime symmetry group of one’s equations of motion is more restricted than the spacetime symmetry group of one’s ontology—for in this case, one would be shirking one’s duty by in effect employing geometrical structure in one’s dynamical theory while not being willing to pay the ontological cost.

It is tempting to think that this picture ought to generalize: Why should spacetime symmetries be special? At the level of slogans, the idea is easy to state—our interpretative practice should be guided by the principle that the symmetries of a theory’s laws and the symmetries of its ontology should coincide. It isn’t obvious how to give a precise and general formulation of the idea. But one thing that seems clear is that any such formulation would have as a consequence the second doctrine discussed in the opening section of this paper:

D2. Two solutions of a classical theory’s equation of motion are related by a symmetry if and only if they are physically equivalent, in the sense

symmetry group of a theory. The situation is more perplexing in the case of fancier quantum theories. On the one hand, an arbitrary C^* -algebra automorphism is pretty clearly the analogue of an arbitrary symplectic or unitary transformation and so is not a good candidate to be a symmetry of a theory: indeed, in many cases of interest any two states are related by such an automorphism; see Kishimoto *et al.* (2003). On the other hand, in some contexts it is not possible to identify symmetries of a theory with those C^* -algebra automorphisms that preserve the Hamiltonian because there *is* no Hamiltonian operator available at the C^* -algebra level; on this point, see Ruetsche (2011, §12.3).

that they are equally well- or ill-suited to represent any particular physical situation.³⁷

We have seen that disaster follows if this doctrine is combined with the notion that the symmetries of a classical theory are the maps that send solutions to solutions. But having set aside that notion, we can ask whether D2 can be safely combined with one of the more nuanced and discriminating notions of symmetry on offer.

Indeed, one can give various plausibility arguments in favour of theses in the neighbourhood of D2.

- (a) Intuitively speaking, two solutions are related by a symmetry if and only if they are interchangeable by the lights of the theory's formalism. And surely if everything has been set up properly (i.e., just the right ingredients have been built into the formalism of the theory), two solutions that are formally interchangeable should also have identical representational capacities—and so should be physically equivalent in the present sense.
- (b) Further, one might think that there is a pretty good reason for thinking that two solutions related by a classical symmetry of a theory must be physically equivalent. For such symmetries can be thought of as smooth transformations of the dependent and independent variables of the theory that preserve the form of the equations of the theory. How could our representational practices reasonably distinguish between two solutions of a theory that are related by a reparameterisation of the theory's variables to which the equations of the theory are themselves indifferent?
- (c) That this is the right way to go is also suggested by consideration of familiar cases. No one denies that solutions of a theory's equations are physically equivalent if they are related by a spacetime symmetry of the theory (i.e., by a symmetry that transforms the independent variables of the theory without affecting the dependent variables).³⁸ A similar

³⁷In truth, in order to formulate a plausible doctrine in the neighbourhood of D2, one should probably work with infinitesimal symmetries—otherwise it is easy to concoct counterexamples by deleting points from a theory's space of solutions.

³⁸There *is* disagreement over the question whether there are distinct physical possibilities related by shifts and the like. But that is a different question.

consensus exists regarding paradigm cases of (global and local) *internal symmetries* (that involve transformations of dependent variables of the theory that do not affect the independent variables). Thus, if one has a theory involving a complex scalar field such that global phase transformations of the form $\phi(x) \mapsto e^{i\theta}\phi(x)$ (θ a constant) are symmetries, then one regards solutions related by such a transformation as physically equivalent. Similarly, one regards as physically equivalent solutions of Maxwell’s equations in vector potential form that are related by local gauge transformations of the form $A^\mu(x) \mapsto A^\mu(x) + d\Lambda(x)$ (here Λ is a smooth function on spacetime, so in general the transformation effected on the dependent variables varies from spacetime point to spacetime point). Why should other symmetries be different?

Unfortunately, these mutually reinforcing half-arguments don’t add up to much: each of the general notions of symmetry that we have considered leads to unpalatable consequences when plugged into D2.

Consider first the notion of a classical symmetry of a differential equation. As noted above, most equations admit relatively few such symmetries, so that solutions related by a classical symmetry share a great deal in common. But there exist equations whose symmetry groups are so large that they act transitively on the space of solutions (i.e., for any two solutions, u_1 and u_2 , there is a classical symmetry of the equation that maps u_1 to u_2). This sort of shocking behaviour can be found in some of the most basic theories of classical physics.

- (i) A single Newtonian free particle. In this case, the spacetime symmetries of the theory are given by the Galilei group, and any solution can be mapped on to any other by such a symmetry.³⁹
- (ii) The harmonic oscillator. In one spatial dimension, an arbitrary solution takes the form $q = A \cos t + B \sin t$. The theory admits a one-parameter group of classical symmetries that acts on solutions by changing the value of B , and another such group that changes the

³⁹In fact, the classical symmetries for this theory are not exhausted by the Galilei symmetries—already in two spacetime dimensions, where the Galilei group is three-dimensional, the classical symmetry group for the free particle is eight-dimensional. See, e.g., Duarte *et al.* (1987).

value of A —so any solution can be mapped to any other by a symmetry of the theory.⁴⁰

- (iii) Linear homogeneous partial differential equations—such as the heat equation, the wave equation, the source-free Maxwell equations, etc. Corresponding to any solution u_0 of such an equation, there is a classical symmetry $T_{u_0} : u \mapsto u + u_0$.⁴¹ And, of course, for any two solutions u_1 and u_2 of a linear equation, there is a solution u_0 such that $u_2 = u_1 + u_0$. So any two solutions of a linear homogeneous partial differential equation are related by a symmetry.

If we maintain that any two solutions related by a classical symmetry are physically equivalent, then each of these theories will be unable to discriminate among the systems it provides good models for—it will consider them to have the same physics. That is a presumably the right verdict in the case of the theory of a free Newtonian particle—if we follow the policy of regarding solutions related by a spacetime symmetry as physically equivalent, then no degrees of freedom remain in this theory once such symmetries are factored out. But the other examples are very different: under any ordinary reading, they admit solutions that represent situations in which nothing is happening (the oscillator is permanently immobile at the origin, the field is in a ground state) and others that represent situations in which plenty is going on (the oscillator is continually sproinging around, energy in the form of heat or waves is propagating). An approach to understanding physical theories that leaves us unable to see these distinctions is not something we can live with. So D2 is false if understood as a thesis concerning classical symmetries of differential equations.

Further, employing generalized or non-local symmetries would not help: the problems just noted would remain, and new ones of the same ilk would crop up.⁴²

⁴⁰See, e.g., Lutzky (1978, §3). Note that the complete classical symmetry group of the one-dimensional oscillator is eight-dimensional, and so outstrips the group of spacetime symmetries. Note also that all of these results carry over, *mutatis mutandis*, to the case of a time-dependent oscillator in n spatial dimensions; see Prince and Eliezer (1980).

⁴¹See, e.g., Hydon (2000, p. 145). For the classical symmetries of the heat equation and the wave equation, see, e.g., Olver (1993, Examples 2.41 and 2.43). For symmetries of the source-free Maxwell equations, see Anco and Pohjanpelto (2008).

⁴²E.g., the group of non-local symmetries of the Kepler problem acts transitively on solutions; see Leach *et al.* (2003).

Nor does it help to shift attention to variational or Hamiltonian symmetries. Consider first variational symmetries. The class of (classical or generalized) variational symmetries of an equation is more restrictive than the class of (classical or generalized) symmetries of an equation. And that has some benefits: the variational symmetries of the wave equation do *not* include addition-of-an-arbitrary-solution (and do not act transitively on the space of solutions).⁴³ But: the problem with the harmonic oscillator remains—for any two solutions, there is a variational symmetry that maps one to the other.⁴⁴ So the class of variational symmetries is still in some respects too generous underwrite D2. It is also in some respects too restrictive: for example, neither the boost symmetry of classical mechanics nor the scaling symmetry of general relativity is a variational symmetry.⁴⁵ But it is hard to deny that two solutions related by such a symmetry are physically equivalent in the relevant sense.⁴⁶

Further, recall that not every equation of motion admits a Lagrangian treatment.⁴⁷ True, with some ingenuity one can often find surrogates for such equations that *do* admit Lagrangian treatment (e.g., by replacing the variables of the original theory by suitable potentials, or by introducing non-physical fields).⁴⁸ But it is hard to see why one should have to take such detours in order to understand the connection between symmetry and

⁴³See Olver (1993, Example 4.15).

⁴⁴See Lutzky (1978) and Prince and Eliezer (1980).

⁴⁵What happens if in place of variational symmetries, we consider divergence symmetries? Some problems go away, others reappear. Good news: Galilean boosts and certain types of scaling transformations count as symmetries (see fn. 33 above). Bad news: addition-of-an-arbitrary solution is a divergence symmetry of the wave equation; Olver (1993, Example 4.36). So if we count solutions related by a divergence symmetry as physically equivalent, we have to view every pair of solutions of the wave equation as begin physically equivalent.

⁴⁶In the case of general relativity, denying that solutions g and $c \cdot g$ (c a positive constant) related by a scale transformation are physically equivalent is especially difficult for those who deny that there are possible worlds that agree about distance ratios but disagree about matters of absolute distance—for according to such philosophers there is only one world that is adequately represented by g and $c \cdot g$ taken together, and it is isn't easy to see how one of the two solutions could have a better claim than the other to represent that world.

⁴⁷The problem of identifying those which do so is known as *the inverse problem of the calculus of variations* or as *Helmholtz's problem*. For an introductory survey, see Prince (2000).

⁴⁸See, e.g., Ibragimov and Kolsrud (2004), Olver (1993, Exercises 5.35, 5.36, and 5.46), Rosen (1966), and Sorkin (2002).

physical equivalence.

What about Hamiltonian symmetries, then? This class is again too restrictive in some respects. As in the variational case, scaling symmetries are typically not Hamiltonian symmetries. Galilean boosts cause the same sort of trouble—since they typically leave the potential energy of a system invariant while altering its kinetic energy, they fail to preserve the Hamiltonian. So if one took solutions to be physically equivalent only if related by a Hamiltonian symmetry, then one would have to violate the general principle that solutions of classical theories are physically equivalent if related by a spacetime symmetry.⁴⁹

At the same time, the class of Hamiltonian symmetries is also too generous in some respects. In the Hamiltonian setting we are at least safe from the threat of the group of symmetries of a theory acting transitively on the space of states (since Hamiltonian symmetries preserve the Hamiltonian and theories ordinarily allow states of differing energy). But we nonetheless still run into cases in which we are unwilling to consider states or solutions related by Hamiltonian symmetries as physically equivalent.

Consider the Kepler problem. This theory has a number of Hamiltonian symmetries, some of which are spacetime symmetries and some of which are not. The spacetime symmetries include time translation (associated with conservation of momentum) and rotation (associated with conservation of angular momentum). As usual, we want to regard solutions related by spacetime symmetries as physically equivalent—in the present case, this means that we regard solutions as physically equivalent if the corresponding elliptical orbits are of the same shape (i.e., have the same eccentricity and have equally long major axes). The further Hamiltonian symmetries of the Kepler problem are associated with the conservation of the Lenz–Runge vector. If two solutions are related by one of these symmetries, then the corresponding ellipses have equally long major axes, but (in general) have different eccentricities and different orientations in space.⁵⁰ The upshot is

⁴⁹Just as one might move from variational symmetries to divergence symmetries (see fn. 33 above), one might consider transformations of a system’s phase space that leave invariant the set of Hamiltonian trajectories without worrying about whether they also leave the Hamiltonian itself invariant. This class would include Galilean boosts. But it would not include scale transformations—e.g., in the Kepler problem, changing the scale used to measure distances carries one from a periodic solution to a periodic solution with a different period, so such a transformation doesn’t correspond to a mere permutation of Hamiltonian trajectories.

⁵⁰See Morehead (2005). This symmetry corresponds to a generalized symmetry of the

that if we take being related by a Hamiltonian symmetry to imply physical equivalence, then we must take solutions of the (negative energy) Kepler problem to be physically equivalent if and only if they correspond to ellipses with equally long major axes.⁵¹ But we do not normally regard highly eccentric orbits and perfectly circular orbits as being physically equivalent.⁵²

We run into the same problem with the harmonic oscillator: the family of Hamiltonian symmetries acts transitively on the surfaces of constant energy in the space of initial data—so two solutions of the theory are related by a Hamiltonian symmetry if and only if they represent the system as having the same total energy.⁵³ But two solutions of the harmonic oscillator can have the same energy while corresponding to quite different motions—e.g., in the two-dimensional case, one might have the particle moving back and forth on a line segment while the other has it executing a circular motion.⁵⁴

equations of motion.

⁵¹For the negative energy Kepler problem, the length of the major axis determines the energy of a solution, so any Hamiltonian symmetry leaves this quantity invariant; see, e.g., Goldstein *et al.* (2002, §3.7).

⁵²(1) One might think that we would do so if scaling transformations were up for grabs, since one can indeed transform a circular orbit into an eccentric orbit by rescaling one coordinate axis while leaving the others invariant. But this sort of rescaling does not preserve the Hamiltonian of the Kepler problem (because it changes the length of the major axis of some solutions). (2) The discussion above glosses over an interesting subtlety. At the infinitesimal level, one does indeed find a large set of symmetries of the Kepler problem. But the space of initial data features a singular set of points (corresponding to situations in which the two particles collide). And the existence of this singular set provides an obstruction to integrating the infinitesimal symmetries into a group action; see Cushman and Bates (1997, p. 74). However: there exist ways of regularizing the singularities of the Kepler problem—and these constructions lead to a family of finite symmetries that act transitively on surfaces of constant energy. For discussion and references, see Cushman and Bates (1997, Chapter II) or Guillemin and Sternberg (1990, §2.7). (3) Note that the fact that the Hamiltonian symmetry group of the Kepler problem is larger than its spacetime symmetry group plays an important role in the quantum theory of the hydrogen atom; see Jauch and Hill (1940) or Guillemin and Sternberg (1990, §7).

⁵³Consider the two-dimensional case. There are four conserved quantities, corresponding to four independent Hamiltonian symmetries. A surface of constant energy is a three-sphere in the phase space and the group of Hamiltonian symmetries is the group $U(2)$, which acts transitively on the energy surfaces. For details, see, e.g., Goldstein *et al.* (2002, §9.8), Cushman and Rod (1982), or Cushman and Bates (1997, Chapter II).

⁵⁴Again, the fact that the symmetry group of the classical system is $U(2)$ rather than just the spacetime symmetry group plays an important role in the quantum theory; see Jauch and Hill (1940).

Or, again, consider the Korteweg–de Vries equation. This equation governs a real-valued field u in one spatial dimension. Under its usual interpretation, this field describes the dynamics of a shallow body of water with one horizontal dimension, with $u(x, t)$ giving the depth of the water above spatial point x at time t . Consider a particular solution, u_0 . How large is the family of solutions physically equivalent to u_0 ? Well, it surely includes solutions related to u_0 by spacetime symmetries, and perhaps also scaling symmetries. But the family of generalized symmetries of the equation and the family of Hamiltonian symmetries are both infinite-dimensional—so it seems clear that each must contain many symmetries that relate physically inequivalent solutions.⁵⁵

7 Outlook

Where does this leave us? There are various interesting formal notions of symmetry applicable to classical physical theories. But none of the standard ones are suited to underwrite a principle like D2 that makes a direct link between the being related by a symmetry and being physically equivalent.

I leave it as a challenge to the reader to identify a general and interesting formal notion of symmetry that renders D2 true—or, better, to identify a family X of symmetries such that two solutions of a theory that are related by a symmetry are physically equivalent if and only if they are related by a symmetry that belongs to X .

⁵⁵In the examples just considered, it is pretty clear that one does not want to count every pair of solutions related by a generalized symmetry as being physically equivalent. Does one *ever* want to count solutions as physically equivalent that are related by a generalized symmetry that is not a classical symmetry? Yes—for instance, when the solutions in question are also related by a respectable classical symmetry. Consider, e.g., the generalized gauge transformations described in fn. 21 above—if two solutions are related by such a symmetry, then they are also related by an ordinary gauge transformation.

Are there pairs of solutions related by a generalized symmetry (but not by any classical symmetry) that one would want to consider physically equivalent? That appears to be a more difficult question. Part of the difficulty lies in the fact that what one has in practice are the infinitesimal generators of generalized symmetries: it is in general a nontrivial task to find the corresponding group actions; see, e.g., Olver (1993, 297 ff.). Further, even in cases where the corresponding groups of transformations can be determined, their physical interpretation can be obscure; see, e.g., Olver (1984, 136 f.).

Above we have seen that we ordinarily take spacetime symmetries to belong to X . Further, we have seen that there are classical symmetries, generalized symmetries, non-local symmetries, variational symmetries, and Hamiltonian symmetries that are not in X . We have also seen that X contains symmetries that fall outside of the classes of spacetime symmetries, variational symmetries, and Hamiltonian symmetries

Perhaps it is possible to find some formal notion of symmetry that combines well with D2. But it is hard to be optimistic—certainly, the ways of encoding the content of laws that are most appealing to mathematicians and physicists appear to lead to notions of symmetry that are coolly indifferent to considerations of representational equivalence. So 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.⁵⁶

Appendix

Here are a few more details about classical symmetries of differential equations. Let's begin by describing $J^k(E)$.⁵⁷ For present purposes, it is easiest to work in terms of coordinates.⁵⁸

A point y in M is specified by specifying a d -tuple of real numbers, (y_1, \dots, y_d) (the coordinates of y).

A point (y, v) in E is specified by specifying a $(d + m)$ -tuple of real numbers, $(y_1, \dots, y_d; v_1, \dots, v_m)$ (the coordinates in M of y plus the coordinates in W of $u(y)$, the value of the field u at spacetime point y).

A point (y, v, p) in $J^k(E)$ is specified by specifying a $\left[d + m \binom{d+k}{k} \right]$ -tuple of real numbers, $(y_1, \dots, y_d; v_1, \dots, v_m; p_1, \dots, p_n)$, where the p_i

⁵⁶Of course, some interesting weaker relative of D2 might be true. E.g., for all that has been said here, being related by a Hamiltonian symmetry that corresponds to a classical symmetry of their equation of motion may be a sufficient condition for two solutions to be physically equivalent.

⁵⁷For what follows, see, e.g., Krasil'shchik and Vinogradov (1999, Chapter 3).

⁵⁸But note that everything described below can be done in a respectable global and coordinate-independent fashion. See, e.g., Saunders (2008).

are just numerous enough to be regarded as a list of the values of the partial derivatives (from order one through k) of u by the spacetime coordinates.

When we think of $J^k(E)$ in this way, it provides a natural way to encode the content of any k^{th} order partial differential equation whose dependent and independent variables parameterise E . For think what such a differential equation does: for each point x of spacetime, the equation imposes a constraint on the values of the field and its partial derivatives through order k when evaluated at x . But, roughly speaking, that is to say that such a differential equation determines a submanifold of $\mathcal{E} \subset J^k(E)$, and that the content of the equation is exhausted by the structure of \mathcal{E} as a submanifold of $J^k(E)$.

Now we are getting somewhere: we can identify symmetries as one-to-one and onto maps from $J^k(E)$ to itself that preserve \mathcal{E} and all relevant structure of the ambient space $J^k(E)$. But what does this structure amount to?

We know that $J^k(E)$ is a manifold and that we are working in a setting in which everything is at least a bit smooth—so it is natural to require our symmetries to be appropriately smooth diffeomorphisms from $J^k(E)$ to itself.

There is one further kind of structure on $J^k(E)$ that matters in the present context. Recall that we have said that, intuitively, the extra variables that we add in moving from E to $J^k(E)$ are supposed to correspond to the values of the partial derivatives of the field with respect to the spacetime coordinates. We need to put some structure on $J^k(E)$ in order to enforce this intuitive demand.

In order to get a feeling for what is required here, consider the following line of thought. Let $u : M \rightarrow W$ be a kinematically possible field. The content of u is encoded in its graph, Γ_u , the subset of E consisting of pairs for the form $(x, u(x))$, for $x \in M$. Of course, the map $x \in M \mapsto (x, u(x)) \in \Gamma_u$ is a diffeomorphism. Further, a sufficiently smooth submanifold $\Gamma \subset E$ corresponds to a kinematically possible u in this way if and only if the projection map $\bar{\pi} : (x, v) \in \Gamma \rightarrow x \in M$ is a diffeomorphism from Γ onto M .

Now, each kinematically possible u also determines a diffeomorphism from M to $J^k(E)$. We define the k -jet of u to be the map $j[u] : x \in M \rightarrow (x, u(x), p) \in J^k(E)$, where $p = (p_1, \dots, p_n)$ encodes the values

at x of the partial derivatives of u through order k . The k -jet of u is a diffeomorphism onto its image J_u . There is of course a projection map $\pi : (x, v, p) \in J^k(E) \mapsto x \in M$, whose restriction to J_u is the inverse of the jet $j[u]$.

It may be tempting to suppose that any sufficiently smooth submanifold $K \subset J^k(E)$ such that $\pi : K \rightarrow M$ is a diffeomorphism onto its image arises as the J_u for some kinematically possible u . But this is false. For consider such a K . For each spacetime point $x \in M$, K includes exactly one point (x, v, p) . So, questions of boundary conditions aside, K does determine a kinematically possible u —for each spacetime point x just set $u(x) = v$, where (x, v, p) is the unique point in K that π sends to x . But there is no guarantee that $J_u = K$, precisely because we have not yet built into the structure of $J^k(E)$ any connection between the p_i and the partial derivatives of the field—so in general, if (x, v, p) is a point in K and u is the kinematically possible field determined by K , there is no reason to expect that the values of the partial derivatives of u at x are given by p (that is, in general one expects that $j[u](x) \neq (x, v, p)$ so that $J_u \neq K$).

There is an elegant solution to this problem. One imposes on $J^k(E)$ some geometric structure, in the form of the *Cartan distribution*, \mathcal{C} , which singles out at each point of $J^k(E)$ a distinguished subspace of dimension $\left[d + m \binom{d+k-1}{d-1} \right]$ in the $\left[d + m \binom{d+k}{k} \right]$ -dimensional tangent space at that point. The Cartan distribution has the following beautiful feature: a submanifold K of $J^k(E)$ that is diffeomorphic to M via the natural projection map $\pi : (x, v, p) \mapsto x$ is of the form J_u for some kinematically possible u if and only if at any point of K , all tangent vectors pointing along K lie in the privileged subspace picked out at that point by \mathcal{C} . That is: the Cartan distribution encodes our intuitive constraint that the p_i should correspond to the values of the partial derivatives of fields.⁵⁹

Putting all of this together: a symmetry (in the present sense) of our equation is a diffeomorphism from $J^k(E)$ to itself that preserves: (i) the submanifold $\mathcal{E} \subset J^k(E)$ that encodes the equation; and (ii) the Cartan distribution \mathcal{C} . It now turns out that (except in the special case where $m = 1$) the only such diffeomorphisms arise via diffeomorphisms from E to itself (this claim will be unpacked in the next paragraph). And of course,

⁵⁹The space of vectors picked out by \mathcal{C} at a point of $J^k(E)$ coincides with the vectors annihilated by the family of one-forms on $J^k(E)$ that enforce the differential relations that require the p_i to be the derivatives of the components of u with respect to the x_j .

since each kinematically possible field u can be identified with the submanifold $j[u](M) \subset J^k(E)$, any diffeomorphism from $J^k(E)$ to itself acts in a natural way on the space of kinematically possible fields—and (subtleties aside) any diffeomorphism from $J^k(E)$ to itself that preserves \mathcal{E} maps solutions to solutions. So: the symmetries (in the present sense) of a differential equation are the transformations of the space of kinematically possible solutions that map solutions to solutions and that are suitably local, in the sense of depending only on the independent and dependent variables of the theory.

There are two unfinished bits of business. The first is to unpack the notion of a diffeomorphism from $J^k(E)$ to itself arising via a diffeomorphism from E to itself. Let $\bar{d} : E \rightarrow E$ be a diffeomorphism. We seek to define a corresponding diffeomorphism $d : J^k(E) \rightarrow J^k(E)$. We proceed by selecting an arbitrary $(x, v, p) \in J^k(E)$ and showing how to find $d(x, v, p)$. Let u be a kinematically possible field such that $u(x) = v$ and $j[u](x) = (x, v, p)$ (i.e., (x, v, p) gives the value of u and its partial derivatives at x). Let Γ_u be the graph of u in E (i.e., the set of points of E of the form $(x, u(x))$, for all $x \in M$). Γ_u is a smooth submanifold of E that projects diffeomorphically to M under the map $\bar{\pi} : (x, v) \mapsto x$. Now, since \bar{d} is a diffeomorphism from E to itself, the set $\bar{d}(\Gamma_u)$ is also a smooth submanifold of E . Sadly, $\bar{d}(\Gamma_u)$ need not project diffeomorphically to M . But let us ignore that detail, and pretend that it does—and hence corresponds to a kinematically possible u^* .⁶⁰ Then we can define $d(x, v, p) := j[u^*](x')$, where $x' = \bar{\pi}(\bar{d}(x, v))$. The resulting map $d : J^k(E) \rightarrow J^k(E)$ is a well-defined diffeomorphism (in particular, it is independent of the choices we made along the way).

The final piece of unfinished business is to note that one typically works with fields that have well-defined transformations laws under changes of coordinates on spacetime: so a spacetime diffeomorphism $\bar{d} : M \rightarrow M$ induces in a natural way a diffeomorphism $\bar{D} : E \rightarrow E$.⁶¹ The corresponding transformation $D : \mathcal{K} \rightarrow \mathcal{K}$ is then given by $D : u(x) \mapsto u(\bar{D}^{-1}(x))$. Then we find, as expected, that D is a classical symmetry if and only if it maps solutions to solutions.

⁶⁰To do everything honestly, we would need to: (i) introduce a generalized notion of a solution; (ii) work locally; or (iii) shift our focus to the infinitesimal symmetries.

⁶¹This is one point at which the story can become more complicated for fields that are not given by tensors. See, e.g., Kolář *et al.* (1993) on gauge natural bundles.

REFERENCES

- Abraham, R. and J. Marsden (1985), *Foundations of Mechanics*. Second edition. Cambridge, MA: Perseus.
- Anco, S. and J. Pohjanpelto (2008), “Generalized Symmetries of Massless Free Fields on Minkowski Space.” *Symmetry, Integrability and Geometry: Methods and Applications* 4: 004.
- Anderson, I. and C. Torre (1996), “Classification of Local Generalized Symmetries for the Vacuum Einstein Equations,” *Communications in Mathematical Physics* 176: 479–539.
- Baker, D. (2010), “Symmetry and the Metaphysics of Physics.” *Philosophy Compass* 5: 1157–1166.
- Bluman, G. (2005), “Connections Between Symmetries and Conservation Laws.” *Symmetry, Integrability and Geometry: Methods and Applications* 1: 011.
- Bluman, G. and S. Anco (2002), *Symmetry and Integration Methods for Differential Equations*. Berlin: Springer–Verlag.
- Brading, K. and E. Castellani (2007), “Symmetries and Invariances in Classical Physics,” in J. Butterfield and J. Earman (eds.), *Philosophy of Physics*. Part B. Amsterdam: Elsevier, 1331–1367.
- Cantwell, B. (2002), *Introduction to Symmetry Analysis*. Cambridge: Cambridge University Press.
- Castrillón López, M. and J. Marsden (2008), “Covariant and Dynamical Reduction for principal Bundle Field Theories,” *Annals of Global Analysis and Geometry* 34: 263–285.

- Cushman, R. and L. Bates (1997), *Global Aspects of Classical Integrable Systems*. Basel: Birkhäuser.
- Cushman, R. and D. Rod (1982), “Reduction of the Semisimple 1:1 Resonance,” *Physica D* 6: 105–112.
- Dasgupta, S. (2010), “Symmetries in Physical Reasoning.” Unpublished manuscript.
- Debs, T. and M. Redhead (2007), *Objectivity, Invariance, and Convention: Symmetry in Physical Science*. Cambridge, MA: Harvard University Press.
- Duarte, L., S. Duarte, and I. Moreira (1987), “One-Dimensional Equations with the Maximum Number of Symmetry Generators,” *Journal of Physics A* 20: L701–L704.
- Earman, J. (1989), *World Enough and Space-Time: Absolute versus Relational Theories of Space and Time*. Cambridge, MA: MIT Press.
- Goldstein, H., C. Poole, and J. Safko (2002), *Classical Mechanics*. Third edition. New York: Addison–Wesley.
- Guillemin, V. and S. Sternberg (1990), *Variations on a Theme by Kepler*. Providence, RI: American Mathematical Society.
- Healey, R. (2009), “Perfect Symmetries,” *British Journal for the Philosophy of Science* 60: 697–720.
- Hydon, P. (2000), *Symmetry Methods and Differential Equations: A Beginner’s Guide*. Cambridge: Cambridge University Press.
- Ibragimov, N. and T. Kolsrud (2004), “Lagrangian Approach to Evolution Equations: Symmetries and Conservation Laws,” *Nonlinear Dynamics* 36: 29–40.
- 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, 371–392.
- Jauch, J. and E. Hill (1940), “On the Problem of Degeneracy in Quantum Mechanics,” *Physical Review* 57: 641–645.
- Kishimoto, A., N. Ozawa, and S. Sakai (2003), “Homogeneity of the Pure State Space of a Separable C^* -Algebra,” *Canadian Mathematical Bulletin* 46: 365–372.

- Klainerman, S. (2008), “Partial Differential Equations,” in T. Gowers, J. Barrow–Green, and I. Leader (eds.), *The Princeton Companion to Mathematics*. Princeton: Princeton University Press, 455–483.
- Kolář, I., P. Michor, and J. Slovák (1993), *Natural Operations in Differential Geometry*. Berlin: Springer–Verlag.
- Krasil’shchik, I. and A. Vinogradov (eds.) (1999), *Symmetries and Conservation Laws for Differential Equations of Mathematical Physics*. Providence, RI: American Mathematical Society.
- Leach, P., K. Andriopoulos, and M. Nucci (2003), “The Ermanno–Bernoulli Constants and the Representations of the Complete Symmetry Group of the Kepler Problem,” *Journal of Mathematical Physics* 44: 4090–4106.
- Lévy–Leblond, J.–M. (1971), “Conservation Laws for Gauge-Variant Lagrangians in Classical Mechanics,” *American Journal of Physics* 39: 502–506.
- Lutzky, M. (1978), “Symmetry Groups and Conserved Quantities for the Harmonic Oscillator,” *Journal of Physics A* 11: 249–258.
- Morehead, J. (2005), “Visualizing the Extra Symmetry of the Kepler Problem,” *American Journal of Physics* 73: 234–239.
- North, J. (2010), “Structure in Classical Mechanics.” Unpublished manuscript.
- Olver, P. (1984), “Conservation Laws in Elasticity. II. Linear Homogeneous Isotropic Elastostatics.” *Archive for Rational Mechanics and Analysis* 85: 131–160.
- Olver, P. (1993), *Applications of Lie Groups to Differential Equations*. Second edition. Berlin: Springer–Verlag.
- Pohjanpelto, J. (1995), “Symmetries, Conservation Laws, and Maxwell’s Equations.” In T. Barrett and D. Grimes (eds.) *Advanced Electromagnetism: Foundations, Theory and Applications*. Singapore: World Scientific, 560–589.
- Prince, G. (2000), “The Inverse Problem in the Calculus of Variations and its Ramifications,” in P. Vassiliou and I. Lisle (eds.) *Geometric Approaches to Differential Equations*. Cambridge: Cambridge University Press, 171–200.
- Prince, G. and C. Eliezer (1980), “Symmetries of the Time-Dependent N -Dimensional Oscillator,” *Journal of Physics A* 13: 815–823.

- Prince, G. and C. Eliezer (1981), “On the Lie Symmetries of the Classical Kepler Problem,” *Journal of Physics A* 14: 587–596.
- Roberts, J. (2008), “A Puzzle about Laws, Symmetries, and Measurability,” *British Journal for the Philosophy of Science*, 59: 143–168.
- Rosen, N. (1966), “Flat Space and Variational Principle,” in B. Hoffmann (ed.) *Perspectives in Geometry and Relativity: Essays in Honor of Václav Hlavatý*. Bloomington: Indiana University Press, 325–327.
- Ruetsche, L. (2011), *Interpreting Quantum Theories*. Oxford University Press.
- Saunders, D. (2008), “Jet Manifolds and Natural Bundles.” In D. Krupka and D. Saunders (eds.) *Handbook of Global Analysis*. Amsterdam: Elsevier, 1035–1068.
- Sorkin, R. (2002), “An Example Relevant to the Kretschmann–Einstein Debate,” *Modern Physics Letters A* 17: 695–700.
- Torre, C. (1995), “Natural Symmetries and Yang–Mills Equations.” *Journal of Mathematical Physics* 36: 2113–2130.