SYMMETRIES IN PHYSICS

Philosophical Reflections

Edited by

KATHERINE BRADING
Wolfson College, University of Oxford

and

ELENA CASTELLANI
University of Florence

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore, São Paulo

Cambridge University Press
The Edinburgh Building, Cambridge CB2 2RU, United Kingdom

Published in the United States of America by Cambridge University Press, New York

www.cambridge.org

Information on this title: www.cambridge.org/9780521821377

First published in print format 2003

ISBN-13 978-0-511-07033-4 eBook (EBL)
ISBN-10 0-511-07033-0 eBook (EBL)

ISBN-13 978-0-521-82137-7 hardback
ISBN-10 0-521-82137-1 hardback

24

Notes on symmetries

GORDON BELOT

These notes discuss some aspects of the sort of symmetry considerations that arise in philosophy of physics. They describe and provide illustrations of: (i) one common sort of symmetry argument; and (ii) a construction that allows one to eliminate symmetries from a given structure.[1] I hope that they suggest a unifying perspective.

1 Symmetries

It is helpful to begin with an abstract characterization of symmetries.

A *structure* consists of: a set, D, of *objects* together with a set, $\mathcal{R} = \{R_i\}_{i \in I}$, of relations defined upon D (no restrictions are placed on the cardinality of D or on that of the index set I). If $(D, \{R_i\}_{i \in I})$ and $(D', \{R'_i\}_{i \in I})$ are structures, then we say that a map $\phi : D \to D'$ *fixes* the n-ary relation R_i if: $R_i(x_1, \ldots, x_n)$ iff $R'_i(\phi(x_1), \ldots, \phi(x_n))$ for every n-tuple of objects in D. The *automorphisms* of the structure $(D, \{R_i\}_{i \in I})$ are bijections $\phi : D \to D$ that fix each $R_i \in \mathcal{R}$. The set of automorphisms forms a group under composition of functions.

There are two approaches to talking about symmetries.[2] Under the first, one identifies the symmetries of a structure with its automorphisms – we will call automorphisms symmetries *in the first sense*. Under the second, symmetries are taken to be permutations of names of objects rather than of the objects themselves. This works as follows. Suppose that we have a first-order language with a predicate symbol, $\mathfrak{R}_i$, for each relation, R_i, of our structure, and enough constants, $a, \ldots, c$, to serve as names for each object in D. A *nomenclature* is a bijection from the set of objects to the set of names. If we fix a nomenclature, then we can consider

[1] Some of this material is developed more fully in Belot (2001; and ‘Dust, time, and symmetry’, unpublished manuscript).

[2] The distinction below is related to that between active and passive symmetries, and coincides with it in some contexts.

the set of atomic sentences true of our structure under that nomenclature – that is, the set of formulae of the form $\mathfrak{R}_i(a, \ldots, c)$ true of our structure under our convention of associating names with objects. Call this set the *complete description* of our structure relative to the given nomenclature. Such a complete description determines the associated structure up to isomorphism. If we now compose the given nomenclature with a permutation of the set of names, we generate a new nomenclature; the permutation is a symmetry *in the second sense* when the complete description relative to the induced nomenclature is identical to that generated by the original nomenclature. The symmetries in the second sense relative to a given nomenclature again form a group under composition, and this group is isomorphic to the group of symmetries in the first sense. Because the two senses are so closely related, it is seldom necessary to decide which sense is in play in a given discussion.

Objects related by a symmetry occupy identical roles in the pattern of relations described by their structure – think of the identity of role of points in Euclidean geometry, or of congruent sides of an isosceles triangle. Below we will be interested in structures whose objects correspond to possibilia – typically, possible objects or worlds. We assume that only appropriately qualitative relations are represented in our structures – so objects related by symmetries will be qualitatively indistinguishable.

2 Symmetry arguments

Symmetry arguments have played a central role in natural philosophical debates, especially those concerning the nature of space, time, and motion. Many of them fall under the following argument form (illustrative examples appear in section 3).

- The point of departure is a given structure, held to provide a representation of the features under investigation that is taken to be (more or less) adequate for the purposes at hand. Very often this structure will be either a representation of (aspects of) a certain spatiotemporal world (in which case it will encode, for example, geometrical facts about the spatiotemporal relations between the parts of the given world) or a space of physical possibilities (which will also typically carry a geometrical structure – though, of course, a non-spatiotemporal one).
- In order to solve some outstanding problem, it is proposed to *extend* the given structure by supplementing its class of relations, yielding an enriched representation of the subject matter.
- We ask whether every symmetry of the original structure is a symmetry of the new structure.
- If not, then we can find a symmetry $\Phi : D \to D$ of the original structure, a new relation R of the extension, and objects $x_1, \ldots, x_n \in D$ such that $R(x_1, \ldots, x_n)$ but not $R(\Phi(x_1), \ldots, \Phi(x_n))$. In this case, the new relations make distinctions between objects which are qualitatively indistinguishable in the original structure.

To the extent that we are confident that the symmetries of the original structure are the 'correct' symmetries at the level at which we are working, the failure of the new relations to respect the symmetries of the original structure provides a reason to reject the proposed new structure – and the problem solution that it serves. The content and force of judgements of correctness will vary from case to case.

- If, on the other hand, one can show that the proposed extension is *invariant*, in the sense that it respects the symmetries of the original structure, then the proposed solution has met a minimum standard. If the extension can be shown to be the unique invariant extension of the sort under consideration, then one has reason to accept the extended structure as a (more or less) adequate representation of the features under investigation – to the extent that one is confident that the type of extension under consideration is indeed the best way to approach the problem at hand.

3 Examples

This section contains five examples of the argument form discussed above, ranging from the ancient to the modern. The first three provide examples in which the structure under investigation is a representation of a single world. In the last two examples the structures are more abstract: in the final example, the structure is the space of worlds possible relative to classical mechanics; in the penultimate example, the structure can be thought of as a coarse-graining of the space of two-particle collision worlds. This selection gives an indication of the range of application of the argument form within natural philosophy.

3.1 Platonic cosmology

The *Timaeus* includes a nice instance of our argument form.

> It is entirely wrong to suppose that there are by nature two opposite regions dividing the universe between them, one 'below,' toward which all things sink that have bodily bulk, the other 'above,' toward which everything is reluctant to rise. For since the whole heaven is spherical in shape, all points which are extreme in virtue of being equally distant from the center must be extremities in just the same manner; while the center, being distant by the same measure from all extremes, must be regarded as the point 'opposite' to them all. . . . When a thing is uniform in every direction, what pair of contrary terms can be applied to it and in what sense could they be properly used? If we further suppose that there is a solid body poised at the center of it all, this body will not move toward any of the points on the extremity, because in every direction they are all alike . . . [3]

Here our initial structure is a highly idealized representation of the Platonic cosmos: a sphere whose distinguished central point represents the Earth.[4] This

[3] 62*c*–63*a*. Translation from Cornford (1997, pp. 262–3).

[4] More precisely: our structure is a spherical subset of Euclidean space, with the usual betweenness and congruence relations defined on the set of points.

structure is invariant under any reflection through a plane through the central point. It follows that it is invariant under rotations (which arise as products of reflections), and hence that any two points on the surface of the sphere are related by a symmetry – and thus count as being qualitatively identical in this representation of the cosmos. Plato supposes that any definition of *down* would have to involve the choice of a distinguished inward-pointing normal to the cosmic sphere. But no such definition – amounting to the supplementation of the original structure by a property possessed only by the single distinguished point – would respect the original symmetries.

Plato moves immediately from this point about definition to a claim about dynamics. In both cases, the claim that the original representation is perspicuous is of course crucial – if we are allowed to bring into play asymmetries in the cosmos or in the central body then the problems Plato considers admit of easy solutions (*up* points towards Polaris; a central shoe moves *toe-wards*).

Here Plato expects his readers to grant as a matter of course that in the sort of investigation he is engaged in, it is only the largest-scale features of the world that should be taken into account. Perhaps to deny this is to mistake the sort of understanding he claims to offer.

3.2 Simultaneity in special relativity

One might run a similar argument to explain to a beginning student Minkowski's claim that with special relativity 'space by itself, and time by itself, are doomed to fade away into mere shadows, and only a kind of union of the two will preserve an independent reality' (1952, p. 75).

In Newtonian spacetime, one has binary relations 'at the same time as' and 'at the same place as' defined on the set of spacetime points. Suppose that we wanted to introduce such structures in Minkowski spacetime.[5] We should, at the very least, require any candidate for 'at the same time as' to be an equivalence relation with three-dimensional, connected, space-like equivalence classes, and require any candidate for 'at the same place as' to be an equivalence relation with one-dimensional, connected, time-like equivalence classes. But there are no such equivalence relations definable on Minkowski spacetime that are also invariant under the symmetries of the spacetime.[6] In this sense, the shift from Newtonian spacetime to Minkowski spacetime deprives us of fully autonomous notions of time and space.

Now, while every student ought to be exposed to this observation, it does little to settle the question of the nature of time in special relativity – for both

[5] Here and below the objects of differential geometry are taken to be structures in which the points of the space are the objects, and the topological and differential structures of the space, along with any tensors defined on it, are taken to be encoded in some appropriate fashion in the relations of the structure.

[6] The only invariant equivalence relations are the trivial ones: the one in which every point is related only to itself, and the one in which every point is related to every other. See Giulini (2001, Theorem 4).

reactionary metaphysicians and heretical philosophers of quantum mechanics deny that Minkowski spacetime provides a complete representation of spatiotemporal reality in the first place. This is quite typical: symmetry arguments are of little polemical value in situations where fundamental questions are at stake, since those are the cases in which there will be little agreement as to whether a given structure provides an acceptable point of departure for such an argument.

But note: in this setting *no one* is interested in using contingent features of the matter distribution to introduce asymmetries which make it (all too) easy to generate invariant simultaneity relations. The question is about the structure of spacetime, and in the context of special relativity it would be cheating to take matter into account.

3.3 Time in dust cosmology

The situation is quite different in general relativistic dust cosmology. There one conceives of spacetime as filled everywhere by dust motes (representing galaxies), which interact with one another only gravitationally – this provides a tractable (and relatively honest) idealization of the large-scale dynamics of the universe for all but the earliest (and, possibly, latest) times (at which other interactions must be taken into account). A solution to the equations consists of a spacetime geometry together with a congruence of time-like geodesics (the worldlines of the dust) and a scalar function (the matter density). The stress-energy of the dust serves as the source term in the Einstein equations.

Now, note that if one augments the structure of Minkowski spacetime by the choice of a privileged congruence of inertial observers at rest relative to one another – in effect imposing a relation 'at the same place as' – then there is a *unique* candidate for 'at the same time as' invariant under the symmetries of the augmented structure.[7] The associated equivalence classes are just the hypersurfaces everywhere orthogonal to the privileged congruence – and coincide with the surfaces of Einstein simultaneity associated with the privileged inertial observers.

Let us now consider a solution of dust cosmology. The dust congruence gives us an analogue of the congruence of freely falling observers in the augmented version of Minkowski spacetime. So it is natural to wonder whether we can define a decent relation 'at the same time as' in our dust cosmology. We look for invariant equivalence relations with connected, three-dimensional equivalence classes each of which intersects each dust worldline exactly once.[8]

[7] Giulini (2001, Theorem 5). See Malament (1977) and Stein (1991) for related results departing from the causal structure of Minkowski spacetime.

[8] The dust congruence and matter density of a dust solution are definable from the metric alone. So in dust cosmology, there is no difference between studying the symmetries of the structure 'spacetime geometry' and studying the symmetries of the structure 'spacetime geometry + material contents'. This equivalence fails for some general relativistic systems, including the Einstein–Maxwell field; see Kramer *et al*. (1980, p. 114).

Such equivalence relations exist in physically realistic dust cosmologies. Indeed, in any dust cosmology with a trivial symmetry group, any partition by space-like hypersurfaces will do – so there will be far too many candidates. But in highly symmetric models, like the Friedmann–Robertson–Walker solutions and the Einstein static universe, there is a *unique* relation of the desired type – the one whose equivalence classes are the hypersurfaces everywhere orthogonal to the dust congruence.

In the early decades of relativistic cosmology, this was taken to suggest that attention to astronomy restored what the local physics of special relativity had dissolved – an objective separation of spacetime into space and time.[9] Gödel was able, however, to construct a dust solution in which there is no invariant equivalence relation possessing the desired features.[10]

On its surface, this shows only that the Minkowski-style argument carries over to Gödel's solutions – a result of some interest, given that a dust congruence is a natural generalization of a structure sufficient to generate an invariant temporal slicing when added to Minkowski spacetime. Gödel himself thought the result had much greater significance.[11]

3.4 Huygens on collision

Here is a bloodless paraphrase of the opening moves in Huygens's analysis of collision (1977, pp. 574–8):

- We consider two bodies A and B. We assume that these bodies are perfectly elastic, of equal mass, and moving along the same line. We assign this line a sense – we call one end 'left' and the other end 'right'. The space of ordered pairs of real numbers is our *space of states*: the ordered pair (a, b) corresponds to a state where A has velocity a along the line while B has velocity b, employing the convention that positive velocities correspond to motion towards the right. This is an impoverished notion of state: we ignore the positions of the bodies on the line.
- We are interested in the dynamics of collision. We seek to define an irreflexive binary relation, $\rightarrow$, on the space of states, where $(a, b) \rightarrow (a', b')$ iff a system initially in state (a, b) eventually evolves into a distinct (a', b').
- We introduce three hypotheses.
  - **Hypothesis I**. The system remains in the initial state unless a collision occurs. So states of the form (a, b) with $a \leq b$ are dead-ends – they do not arrow anything.
  - **Hypothesis II**. For $a > 0$, $(a, -a) \rightarrow (-a, a)$.

[9] See Eddington (1920, p. 163) and Jeans (1936, p. 21ff.).

[10] Gödel (1949a, p. 447); (1949b, p. 560). It is possible to construct hypersurfaces of orthogonality only when the solution is non-rotating. But rotation alone does not suffice to rule out a relation of the desired sort – in Gödel's later, somewhat more realistic, expanding rotating solutions the surfaces of constant matter density have the properties required above of instants; see Gödel (1952).

[11] See Belot, 'Dust, time, and symmetry', unpublished manuscript.

- **Hypothesis III**. For any real number x, if $(a, b) \rightarrow (a', b')$ then $(a + x, b + x) \rightarrow (a' + x, b' + x)$.

- These hypotheses suffice to determine $\rightarrow$. In light of Hypothesis I, we need only investigate states (a, b) with $a > b$.
  - **Proposition I.** For $a > 0$, $(a, 0) \rightarrow (0, a)$, while for $b < 0$, $(0, b) \rightarrow (b, 0)$.
    Proof. Consider the first case. Let $x = -a/2$. Then $(a + x, 0 + x) = (a/2, -a/2)$. By Hypothesis II, $(a/2, -a/2) \rightarrow (-a/2, a/2)$. So by Hypothesis III, $(a/2 - x, -a/2 - x) \rightarrow (-a/2 - x, a/2 - x)$. That is, $(a, 0) \rightarrow (0, a)$. □
  - **Proposition II.** For $a > b$, $(a, b) \rightarrow (b, a)$.
    Proof. As above, choosing this time $x = -\frac{1}{2}(a + b)$. □

This analysis can be recast in our canonical form. Take as the set of objects our space of states and equip it with a binary relation, $\sim$, such that $(a, b) \sim (c, d)$ iff there is a real number x such that $(c, d) = (a + x, b + x)$.[12] Take as our problem the construction of a dynamics, to be encoded in an arrow relation satisfying Hypotheses I and II. Huygens shows that there is unique such relation that respects the symmetries of the original structure.

Hypothesis III, the principle of relativity of inertial motion, is the lynchpin of the analysis – and the element most likely to be challenged by Huygens' Cartesian contemporaries.[13] The success of this principle in solving the problem of collision, and others, ought to be the chief reason given for its acceptance.

In Huygens' proofs, the colliding objects are passed at the moment of collision from the hands of a sailor on a boat gliding down a river to the hands of a confederate at rest on the bank; the situation is arranged so that the solution of the problem for one party follows from Hypothesis II; the other party is then able to solve the problem by appeal to Hypothesis III. This suggests that Huygens viewed each dynamical state as representing the velocities of the colliding bodies relative to some observer, and viewed states related by boosts as descriptions of the *same* system from the point of view of different observers. From this perspective, the principle of relativity says that there is a single set of rules for predicting the outcome of collisions given the initial states, and that observers obtain the correct result by applying this rule to their own description of the initial state.

But there is a second way of interpreting this little theory – one more appealing, probably, to modern eyes than it would have been to Huygens.[14] We can interpret the

[12] This enforces Hypothesis III by ensuring that the symmetries of the original structure are the maps of the form $(a, b) \mapsto (a + x, b + x)$.

[13] Descartes employs the same basic framework in his laws of motion and collision (1991, Part II, sections 36–53) – but Huygens' Propositions I and II directly contradict Descartes' third and sixth rules of impact, while Huygens' Hypothesis I follows from Descartes' first and second laws of motion, and Huygens' Hypothesis II appears as Descartes' first rule of impact. So a Cartesian interested in upholding Descartes' analysis of impact against that of Huygens must locate the error of the latter in the principle of relativity (a principle that plays no role in Cartesian physics).

[14] Of course, some contemporary commentators prefer the original approach; see for example Brown and Sypel (1995). See also note 19, below.

states of the theory as encoding instantaneous absolute velocities of the particles existing in two-particle worlds. States related by boosts now represent distinct possibilities rather than distinct descriptions; and the principle of relativity tells us about the relation between the dynamics of collisions at sets of worlds related by boosts. Note that states represent sets of worlds rather than individual worlds, since they encode no information about the location of the particles – there will be many worlds corresponding to a state of the form (a, a).

3.5 Symmetries in classical mechanics

Suppose that we have n particles of equal mass moving in Euclidean space, subject to (inter-particle and/or external) forces that depend on the location of the particles in space. Then we can cast our dynamics in Hamiltonian form, writing a state of the system as (q, p) where the $3n$-vector q encodes the position coordinates of the particles while the $3n$-vector p encodes their momentum coordinates. We equip the *phase space* $T^*Q := \{(q, p)\}$ with a tensor, the *symplectic form*, $\omega := \sum dq^i \wedge dp_i$, and a scalar function, the *Hamiltonian*, $H := \frac{1}{2}|p|^2 + V(q)$.[15] The first term in the Hamiltonian is the kinetic energy, the second is the potential energy; the components of the forces on the particles are given by $-\frac{\partial V}{\partial q^i}$. There is a unique vector field, X_H, on T^*Q such that $\omega(X_H, \cdot) = dH$.[16] The flow generated by this vector field gives the dynamics of the theory. We write $(q, p) \to_t (q', p')$ when the state (q, p) evolves into the state (q', p') after t units of time.

We can study the symmetries of the structure 'phase space + symplectic form + Hamiltonian'. Because we employ smooth objects, a symmetry will be a diffeomorphism from the phase space to itself that preserves the symplectic form and the Hamiltonian.[17] Because a symmetry preserves the structures that determine the dynamics, the dynamics is also invariant under symmetries – that is, if $\phi : T^*Q \to T^*Q$ is a symmetry and $(q, p) \to_t (q', p')$, then $\phi(q, p) \to_t \phi(q', p')$.

If we restrict attention to forces, such as gravity, that depend only on the inter-particle distances (and not on the location of particles in absolute space) then transformations that correspond to shifting the system in Euclidean space,

[15] The symplectic form is a closed non-degenerate 2-form. The variety of dynamics defined below can be constructed whenever one has a manifold equipped with such a form.

[16] Note that the symplectic form on T^*Q induces the *Poisson bracket* – a Lie bracket satisfying Leibniz's rule – on the space of smooth function on T^*Q via the rule: $\{f, g\} := \omega(X_f, X_g)$. We can write the dynamics in terms of this bracket as $\dot{f} = \{f, H\}$. The advantage of this form is that it applies whenever we have a *Poisson manifold* – a manifold whose space of functions is equipped with a Lie bracket satisfying Leibniz's rule. The notion of a Poisson manifold is more general than that of a symplectic manifold – indeed, every Poisson manifold can be decomposed as a disjoint union of symplectic manifolds.

[17] The group of diffeomorphisms which preserve the symplectic form is immense – it is infinite-dimensional. But the symmetry group of the full structure will be much smaller – at most, of dimension $3n$ (in the case of an integrable system).

or reorienting it by a rotation, will be symmetries.[18] The fact that the operation of shifting the entire system in Euclidean space is a symmetry shows that there is no dynamically preferred origin in the Euclidean space the system inhabits – for the dynamical behaviour of a system whose centre of mass was at the hypothetical origin would be indistinguishable from that of an otherwise similar system that had been shifted (i.e. there is no invariant way to privilege a set of points in the phase space as representing the system as being located at the spatial origin).

So far we have been thinking of the points in our structure as representing possible complete dynamical instantaneous states of a system of particles. But given the determinism (modulo certain technicalities that we can ignore here) of classical mechanics, we might just as well think of them as representing complete possible physical histories of the particles – the specification of a state at a given time is enough to determine the entire history of the system. Thus we can take our phase space to be a space of physically possible worlds, carrying a geometrical structure and scalar which determine the dynamics – where now if $(q, p) \rightarrow_t (q', p')$ then the worlds (q', p') and (q, p) have the same sequence of instantaneous states, with these states occurring t units of time later in one world than in the other.[19]

4 Symmetries of solutions and symmetries of laws

These examples are pretty typical of those one comes across in philosophy of physics – most of the structures that arise in the course of symmetry arguments are either representations (generally, highly idealized ones) of a given spatiotemporal world, or spaces of such representations carrying a structure that encodes the dynamics of a physical theory.

Let's put it this way: suppose that we are interested in philosophical aspects of some physical theory; then we will spend some time contemplating the physics of individual solutions of the equations of the theory, and some time contemplating the content of the laws of the theory by studying (what physicists and mathematicians have to say about) the space of solutions to the equations.

I will make a few remarks about these two occupations and the relation between them.

[18] Galileian boosts are not symmetries in the present sense – they preserve the dynamical trajectories, but not the Hamiltonian.

[19] The statement in the text is tendentious – it is a matter of controversy whether two worlds can differ in this way. My view is that if one denies that the application of time translation (or any other symmetry) generates distinct physical possibilities, then one ought to prefer to the standard formulations of classical mechanics those in which the offending symmetry has been factored out. See section 5 below and Belot (2001).

4.1 Structuring the space of solutions

In the previous section we saw a couple of examples where we were *handed* structures whose objects represented (complete or partial) characterizations of possible worlds, and which came equipped with enough structure to single out a physically interesting group of symmetries. The case of classical particle mechanics, section 3.5, is entirely typical: the symmetries of classical physical theories are studied by studying the symmetries of joint structure composed of the space of states/solutions of the theory and the dynamics-determining structures defined on that space.[20]

Where do these latter structures come from?

This isn't a frivolous question – for if we just look at the differential equations of a theory or at the corresponding space of solutions, we will not have enough structure to pick out the physically interesting symmetries. If we take the latter route, and study the set of solutions as an unstructured set, then arbitrary permutations of the set will count as symmetries. If we base our analysis instead on the differential equations of the theory, we will get a bit further: the space of solutions will plausibly be equipped with a topological and a differential structure, arising out of the use of continuous variables; so the symmetries will at least be diffeomorphisms on the space of states. In either case, it will be possible to relate any pair of solutions by a symmetry – a solution of Newton's gravitational equations in which the planets are all falling into the Sun will count as 'equivalent' to one in which they are in stable orbits. This is a disaster.

Furthermore, the interesting connection between symmetries and conservation laws will be absent under such approaches. This connection can be established when the differential equations of the theory arise as the equations of motion for a Lagrangian or Hamiltonian formulation. In this case, the space of solutions is equipped with a natural geometric structure, the symmetries of the joint structure 'space of states + geometric structure + Lagrangian or Hamiltonian' has as its symmetries the intuitively correct symmetries of the theory, and each (continuous) symmetry of this structure is associated with a conservation law for the original equation(s) of the theory.[21]

[20] The spaces of states of quantum theories also come equipped with familiar structures defining the corresponding dynamics. For reasons of convenience I discuss only the classical case here. But it is interesting (and perhaps important) to note that there is a sense in which quantum theories are special cases of classical theories, as dynamical structures. (i) The space of rays of a Hilbert space carries a symplectic structure; the Schrödinger dynamics are given by solving for the vector field associated with the function on this symplectic manifold given by the expectation value of the Hamiltonian operator; see Landsman (1998, section I.2.5) or Ashtekar and Schilling (1999). (ii) More generally, the space of states of a C^*-algebra carries a dynamically relevant Poisson structure; see Landsman (1998, Proposition I.2.6.8 and Theorem I.3.8.1).

[21] If we set out from a Lagrangian formalism, then there is a natural way to equip the space of solutions of the associated Euler–Lagrange equations with a closed 2-form; this form will be non-degenerate (and hence symplectic) if the Euler–Lagrange equations have a well-posed initial value problem. See Deligne and Freed (1999, section 2).

That the setting of the theories of classical physics within the unifying framework of geometrical mechanics has proved so immensely fruitful suggests that we ought to take as basic, not the equations of motion of classical physics, but rather the Lagrangian or Hamiltonian formulations that give rise to them.[22] This is the implicit or explicit practice of most mathematicians and physicists. But what exactly this 'taking as basic' commits us to is a philosophical question that remains largely unexplored.[23] At first blush it is tempting to think that we are being told to take the set of worlds physically possible relative to some theory as carrying a structure determined by the physics itself, and that this ought to have some consequences for philosophical debates about, for example, the nature of laws, or for the nature of physical possibility more generally. But making out these claims is not a straightforward matter.

4.2 Symmetries of equations vs. symmetries of solutions

- In the case of Newtonian gravitating point particles we find, of course, that the symmetries of the laws – translation in time and the Euclidean symmetries – are not symmetries of typical solutions: only zero-particle solutions are invariant under translations in Euclidean space, and only static solutions are invariant under time translation. This is, in fact, a general feature of differential equations: generic solutions have less symmetry than do the equations that determine them. See Olver (1993, chapter 3).
- Some questions, like that of the existence of a preferred parity, admit of two construals – one focusing on properties of solutions, the other on properties of the laws. Thus we can ask whether there are types of objects or properties such that nature prefers one parity over the other (a question we can reconstrue as whether an imbalance between parities appears in solutions which are good representations of our world); or we can ask whether the laws are invariant under reversal of parity. And, of course, the latter question can be given an affirmative answer even if there are systematic preferences for one parity over the other in nature – although preferences in certain fundamental cases will suggest asymmetries in the laws.

If we set out from a strictly Hamiltonian formalism, then we begin with a space of states/space of initial data equipped with a symplectic form, which we can as usual pull back to the space of solutions by the isomorphism between the spaces induced by fixing a time at which the initial data are posed. More generally, we might work with a space of states equipped with a presymplectic form (a closed but possibly degenerate 2-form whose foliation by null manifolds has a well-behaved leaf-space) – in which case the usual construction equips the space of solutions with a similar form.

22 The structures on the space of solutions discussed in the previous paragraph need not be the end of the story – for example, often we will want to view this space as a cotangent bundle over a configuration space, and view the kinetic term in the Lagrangian or Hamiltonian as arising from a Riemannian metric on the configuration space.

23 See, however, J. Butterfield, 'Solving all problems, postulating all states: some philosophical morals of analytic mechanics', unpublished manuscript.

Now, in the case of parity, the laws version of the question is taken to be much deeper than the particular-solution version of the question. Why should that be? These days, physicists are acutely aware that their present theories are, strictly speaking, false – and much of the most creative and influential work in physics is directed towards creating new theories. The result is that the question 'What do present theories tell us about the world?' becomes 'What hints do present theories contain about future physics?' And while the physics of particular solutions undoubtedly looms large in answers to the first question, it is swamped by considerations relating to structural features of the laws in approaches to the second question. (For further discussion of parity, see Pooley, this volume.)

- Curie's principle is often taken to forbid the evolution of a system from a symmetric state into an asymmetric state. (For the original formulation, see Curie, this volume.)

The principle is true for a large class of theories.[24] Suppose that we are given a space, X, of dynamical states, along with a deterministic dynamics – i.e. for each $x \in X$ and $t \in \mathbb{R}$, there is a unique x' that x evolves into after t units of time; we write $x \rightarrow_t x'$. Suppose, further, that if a state counts as symmetric, it is in virtue of being left invariant by a non-trivial, physically relevant transformation $\Phi : X \rightarrow X$. Finally, suppose that every Φ arising in this way is a symmetry of the dynamics, in the sense that $x \rightarrow_t x'$ implies $\Phi(x) \rightarrow_t \Phi(x')$. Then if x is a symmetric state, so is each x' such that $x \rightarrow_t x'$: by the symmetry of the dynamics, $\Phi(x) \rightarrow_t \Phi(x')$; and by the symmetry of x, $\Phi(x) = x$; so $x \rightarrow_t x'$ and $x \rightarrow_t \Phi(x')$; so by the determinism of the dynamics, $x' = \Phi(x')$.

So in the classical realm, Curie's principle holds so long as the symmetry operations performable on states are also symmetries of the dynamics. This appears to hold true for realistic systems. But it is easy enough to violate for artificial examples.

Consider a Hamiltonian theory of three point particles. Let us say that a state is *equilateral* if it represents the particles as forming an equilateral triangle, with momenta of equal magnitude directed in the same sense along the angle bisectors. Let us say that a state is *scalene* if it represents the particles as forming a scalene triangle. Let (q, p) be some equilateral state; and let Φ be a transformation on the phase space corresponding to the action of reflection in Euclidean space through one of the angle bisectors. Φ permutes the states of our theory, but leaves (q, p) invariant – and it is in virtue of being invariant under such Φ that an equilateral state counts as symmetric. If gravity is the only force acting on our particles, then each such Φ, being a Euclidean symmetry, is a symmetry of our dynamics – and hence leaves (q', p') invariant if $(q, p) \rightarrow_t (q', p')$, so that Curie's principle is satisfied. But we can also consider deviant theories,

[24] See Ismael (1997); J. Earman, 'Spontaneous symmetry breaking for philosopers', unpublished manuscript, and the references therein.

constructed by messily altering the expression for the gravitational potential energy outside of some open neighbourhood of (q, p), so that Φ is no longer a symmetry of the dynamics. In such a theory we have $(q, p) \rightarrow_t (q'', p'')$ with, in general, $(q'', p'') \neq (q', p')$ and (q'', p'') a scalene state. So we have a violation of Curie's principle – the initial state is symmetric with respect to the configuration variables, momenta, and forces, but it evolves into an asymmetric configuration.

Enthusiasts of the principle may be tempted to dismiss such counter-examples. After all, our equilateral state evolves into a scalene state – so one of the three originally congruent sides of the triangle ends up being the longest. In the simplest cases, this will be because the potential of the new theory encodes forces which differentiate between directions in Euclidean space, or treat, say, the first particle differently from the second and the third. It is tempting, perhaps, to think that more complex cases are just variants on these two options. And if this is so, then the example under consideration ought not to be viewed as a counter-example to Curie's principle, for the initial equilateral state is not genuinely symmetric – reflection or interchange of particles is not a *real* symmetry of the initial state, given the sort of information that the potential function takes into account.

Now there is something funny about this objection, since by construction the potential is identical to the Newtonian gravitational potential on an open neighbourhood of the initial equilateral state. So, in effect, we are told the initial state is not symmetric because ... it later evolves to an asymmetric state.

In any case, the presupposition of the objection – that asymmetries in our final state must be grounded in the potential's caring about directions in space or the identity of the particles – is mistaken. We could have first eliminated Euclidean and permutation symmetries from the Newtonian theory (see section 5.2 below) before perturbing the potential. The result would have been a theory in which the particle states are characterized by relative distances and velocities (so the particles live in a relational space, without absolute directions), and in which the state of the system is characterized by a set of three-particle states rather than by an ordered triple of particle states (so that there is no longer a question of which particle is which). Nonetheless, it is possible to rig the potential so that an initial state in which the three relative distances are equal and the three relative velocities are equal evolves into a state in which the particles form a scalene triangle. Only great stubbornness could lead someone to insist that such an initial state should not count as symmetric.

5 Quotienting out symmetries

For every structure $\mathcal{S} = (D, \{R_i\}_{i \in I})$, we can define the associated *quotient structure* $\bar{\mathcal{S}}$, which arises by factoring out the symmetries of $\mathcal{S}$. We define an equivalence relation, $\sim$, on D by declaring $x \sim y$ whenever there exists a symmetry of $\mathcal{S}$,

$\phi : D \to D$ with $y = \phi(x)$. The equivalence class of x under this relation is denoted $[x] := \{y \in D : x \sim y\}$. $\bar{\mathcal{S}}$ has as its set of objects the set $[D] := \{[x] : x \in D\}$, of equivalence classes of $\sim$. For each n-ary relation R_i of $\mathcal{S}$, $\bar{\mathcal{S}}$ has an n-ary relation $[R_i] := \{([x_1], \ldots, [x_n]) : (x_1, \ldots, x_n) \in R_i\}$.

- If $\mathcal{S}$ admits no non-trivial symmetries, then $\bar{\mathcal{S}}$ and $\mathcal{S}$ are isomorphic. If $\mathcal{S}$ is homogeneous – if every pair of its objects are related by a symmetry, as in the case of the order structure of the integers or the rationals – then the quotient has a single object, related to itself by the counterparts of all the non-vacuous relations of $\mathcal{S}$. The interesting cases lie in the intermediate region.
- We can use the same language to describe both structures (using the same predicate symbol, '$\mathcal{R}_i$', as a name for both R_i and $[R_i]$). Choosing names for our new objects generates a complete description of the quotient structure.

  The complete descriptions of the two structures will be closely related. A complete description of the original structure can be transformed into a complete description of the quotient structure by taking names of objects related by symmetries to name identical objects.

  We can also consider the relations between theories describing the two structures. For instance, any constant-free sentence that employs only one-place predicates or that is free of negation symbols will be true of the quotient if it is true of the original structure.[25] But sentences combining multi-place predicates and negation symbols need not be true: let $\mathcal{S}$ be the countable structure whose sole relation, R, gives it the order structure of the integers; then $\bar{\mathcal{S}}$ has a single object, $[x]$, with $[R]([x], [x])$. So the sentence $\forall x \sim \mathcal{R}(x, x)$ is true in $\mathcal{S}$ but false in $\bar{\mathcal{S}}$.
- More generally, we can quotient out by the action of a subgroup of the full symmetry group – declare two objects to be equivalent if related by an element of the chosen subgroup, go on to take equivalence classes, etc.

5.1 Quotienting solutions

If one has a description of a possible world – in the form of a solution of a differential equation, for instance – one can ask whether it admits any symmetries. If it does, then one can consider the distinct description that arises as the quotient of the original.

Advocates of the Principle of the Identity of Indiscernibles will want to deny that any description admitting symmetries corresponds directly to a possible world, while granting that the related quotient (*ceteris paribus*) does so.[26] Indeed, they can go on to insist that a description admitting symmetries is merely a misdescription

[25] Indeed, something stronger is true: if a constant-free sentence is in negation normal form (so that any negation symbols apply to atomic formulae) and each of its negation symbols applies to a one-place predicate, then its truth in the original structure implies its truth in the quotient structure. This follows from Hodges (1997), Theorem 8.3.3(a), since the map $x \mapsto [x]$ is a surjective homomorphism that fixes one-place relations.

[26] For this attitude, see Hacking (1975) and L. Smolin, 'The present moment in quantum cosmology: challenges to the arguments for the elimination of time', PITT-PHIL-SCI 00000153.

of the corresponding quotient, under the strange convention according to which some objects are given multiple names.

To most, this will appear unmotivated at best.[27] Like Black (this volume), many suppose that there could be a Euclidean world, otherwise empty except for two identically constituted iron spheres. What motivation is there for saying that the only possibility in the neighbourhood is the uglier world, containing a single such sphere, related geometrically to itself in its strange non-Euclidean space? Even in cases where there is no gain in awkwardness in passing from the symmetric description to its quotient – as in passing from the covering spacetime of a non-simply connected spacetime to the non-simply connected spacetime itself – many will still feel that there are two genuine possibilities in the neighbourhood.

5.2 Quotienting the space of solutions

If we take the space of solutions of a classical physical theory to be equipped with the sort of rich structure discussed in section 4.1 above, then taking the quotient by the action of a group of symmetries often leads to an interesting result – sometimes with interpretative implications. There exists a large mathematical literature on this technique.[28] I think that the techniques and results of this literature promise to offer a unifying perspective on a number of classic problems in philosophy of physics (the relation between the nature of space and the nature of motion in Newtonian physics, identical particles, the nature and significance of gauge freedom and general covariance).

Let us return to our example of n gravitating point particles.

- Example: Euclidean symmetries. Consider the theory of n gravitating Newtonian point particles. The symmetries of Euclidean space – translation, rotations, reflections, and their products – are symmetries of this theory (action on initial data by one of these symmetries transforms the dynamical trajectory by the action of same symmetry). Call this group $E(3)$. Translations and rotations are generators of the continuous symmetries of the theory, and the corresponding conserved quantities are the total linear and (centre of mass) angular momentum.

  Let δ be the set of points which represent the particles as forming symmetric configurations in Euclidean space. Let Δ be the set of collision points of the phase space, representing states in which two or more particles occupy the same point of Euclidean space. These are sets of measure zero which we excise from the phase space.[29] Call the resulting phase space M. We are also interested in

[27] Hacking's motivation, I take it, stems from his views on the nature of logic; see Hacking (1978; 1979).
[28] For introductions to the mathematics, see Marsden (1992), Marsden and Ratiu (1994), and Singer (2001).
[29] The excision of δ is a convenience that allows us to sidestep complications in the construction of the quotient theory; see Belot (2003, section 10) for discussion and references. The excision of Δ is more essential, since certain types of collision singularity are intractable.

a second space, M_0, the subspace of M corresponding to states in which the system has vanishing linear and angular momentum. Since M_0 is a dynamically closed subspace (being defined by the vanishing of conserved quantities), we can consider the dynamical theory defined upon it by structure inherited from M. We are interested in the quotient structures, $\bar{M}$ and $\bar{M}_0$, that arise when we take the quotient of these spaces by the action of $E(3)$. Each of these encodes a mathematically well-behaved physical theory.[30]

$\bar{M}_0$ gives a theory closely related to the Barbour–Bertotti form of relational dynamics (Barbour and Bertotti, 1982). The points of the quotient space are parameterized by the relative distances and relative velocities of the particles. The dynamics is such that specifying an initial point determines a dynamical trajectory – which gives the same evolution of the relative distances and relative velocities as one gets if one chooses an initial point in the full Newtonian theory with the same pattern of relative distances and relative velocities for the particles and with vanishing angular momentum, then reads off the subsequent values of the relative distances and relative velocities from the Newtonian evolution. Thus $\bar{M}_0$ is an attractive relational theory of motion: it is, mathematically, of the same form as the Newtonian theory (one does not, for instance, have to specify higher derivatives in order to get a well-posed initial value problem); and the relative distances and relative velocities between the particles form a dynamically closed set, whose evolution is deterministic; furthermore the predictions for the empirically accessible variables match those of the Newtonian theory on the cosmologically relevant non-rotating sector. It is also a theory which is naturally set in the space of a classical relationalist about Euclidean space: the theory recognizes no difference between the state of a system located *here*, and one differing only by being shifted over *there*.

The story with $\bar{M}$ is a bit more complicated. We started in M with $6n$ variables (corresponding to the components of the position and momentum of each particle), then we eliminated six of these by identifying points related by the six-dimensional group of Euclidean symmetries. $3n - 6$ of the remaining variables fix the relative distances of the particles and $3n - 6$ of them fix the relative velocities between the particles. Three further variables correspond to the velocity of the centre of mass of the system – and can simply be dropped from the theory, since they are dynamically inert. Informally speaking, the final three variables encode information about the angular momentum of the system – these variables stand in the way of the most straightforward sort of relativist/relationalist interpretation.[31]

[30] Each inherits a Hamiltonian from M, and carries a geometric structure adequate to determine a dynamics – $\bar{M}$ is a Poisson manifold while $\bar{M}_0$ is a symplectic manifold.

[31] These variables correspond to the components of the angular momentum in a frame rotating with the system (rather than in the usual spatially fixed frame). The magnitude of the corresponding vector is preserved, but its direction evolves in time.

- Example: permutation symmetries. Let's alter the notation slightly: now we write the states as $((\vec{q}_1, \vec{p}_1), \ldots, (\vec{q}_n, \vec{p}_n))$ where $\vec{q}_i$ is the position vector of the ith particle and $\vec{p}_i$ is the corresponding momentum. Now consider the transformation $\pi_{(12)} : ((\vec{q}_1, \vec{p}_1), \ldots, (\vec{q}_n, \vec{p}_n)) \mapsto ((\vec{q}_2, \vec{p}_2), (\vec{q}_1, \vec{p}_1), (\vec{q}_3, \vec{p}_3), \ldots, (\vec{q}_n, \vec{p}_n))$, and the corresponding transformations $\{\pi_{(ij)}\}$ for $1 \leq i < j \leq n$. Each of these is a symmetry of our theory, so the group that they generate, S_n, is a symmetry group of the theory.[32] This group permutes the identity of the particles (see French and Rickles, this volume, for further discussion). Let us denote by $\hat{M}$ the quotient of M by the action of S_n. This is again a well-behaved physical theory.[33] We can think of its states as consisting of sets $\{(\vec{q}_1, \vec{p}_1), \ldots, (\vec{q}_n, \vec{p}_n)\}$ of particle states rather than ordered tuples of particle states – the theory keeps track of how many particles are in each state, rather than which particle is in which state. It is an interesting fact that M and $\hat{M}$ underwrite the same statistical theories (see ter Haar, 1995, section 5.9; Huggett, 1999). The reason is as follows. In classical statistical mechanics, one is interested in the ratio of phase space volumes. Now the volume measures in play in M and its subspaces are S_n-invariant, deriving ultimately from the S_n-invariant symplectic form. And the physically interesting subsets of M are also S_n-invariant. The volume of the image of such a subset in $\hat{M}$ will just be $\frac{1}{n!}$ of its volume in M ($n!$ counts the number of elements in S_n). So we will get the same answer whether we measure the ratio of two such regions in M, or measure the ratio of their images in $\hat{M}$.

  Now, the difference between Maxwell–Boltzmann statistics and Bose–Einstein statistics is often informally cashed out in terms of ways of counting possible outcomes of double coin-flips: under the first set of statistics, there are two ways to get {H,T}, under the second, only one. Does the argument above show that classical particles obey Bose–Einstein statistics after all? *No* (see Huggett, 1999, pp. 16–17). A better way of describing the difference between the two sorts of statistics is to say that under one {H,T} is twice as likely as either {H,H} or {T,T}, while under the other these three alternatives are equiprobable. But recall that in constructing M we eliminated all points in which there were particle collisions, including all states in which distinct particles share the same position and momentum – so we cannot make the comparison of the likelihood of classes of states fixed by permutations with classes of states not fixed by permutations in M, or the corresponding comparisons in $\bar{M}$. [34]

[32] Note that while it was a mere notational convenience to take the particles to be of equal mass in previous examples, it is here a necessity.

[33] It is a symplectic manifold, equipped with the projected Hamiltonian.

[34] What happens if we restore to our phase space the subset of Δ consisting of the physically bizarre states fixed by elements of S_n, in which two or more particles share the same position and momentum? On the one hand, not much: this set and its image under the projection to $\hat{M}$ are of measure zero in their respective spaces – so we don't get any difference between the relative weight assigned to these two sets in their respective spaces. On the other hand, there is the following suggestive fact: for a generic (permutation symmetry-free) state there

Whenever we compare two spaces of possible worlds, one the quotient of the other, we are contrasting two ways of counting possibilities. In the two examples discussed above, and in other examples that arise naturally in philosophy of physics, we are in essence faced with the difference between a relatively haecceitistic means of counting possibilities and a relatively anti-haecceitistic means of counting possibilities.[35] In M, but not in $\bar{M}$, if it is possible for the particles to be *thus and so* and to be *here*, then it is also possible for them to be *thus and so* and to be *there* – with the two states qualitatively identical, and differing only as to which spacetime points are occupied. In M, but not in $\hat{M}$, if it is possible for the particles to be *thus and so* with *this* particle playing a certain role, then it is also possible for them to be *thus and so* with *that* particle playing the given role – with the two states qualitatively identical, and differing only in their distribution of roles to particles.

Note that while in the case of the Euclidean symmetries, it was natural to think of (some versions of) the quotienting procedure as leading to the elimination of spacetime points, there is no such straightforward ontological purge in the identical-particles case. For even after the quotient has been taken, the space of states is rich enough to associate with each dynamically possible history a set of continuous spacetime trajectories, labelled by features such as mass, etc. – so there is little impetus to say that anything has been eliminated from the ontology of the worlds described by the theory. Why the difference? Well, the existence of spacetime points is closely tied up with questions of counting of possibilities – so they are vulnerable to elimination in the transition from a haecceitistic means of counting to an anti-haecceitistic one. But in the case of particles we have much more to hang on to.

There is also a crucial technical distinction between the two cases. Taking the quotient of a phase space by the action a discrete group of symmetries yields a phase space that is almost everywhere locally equivalent to the original. This is not so if one works with a continuous group: the presence of continuous symmetries indicates the inclusion of dynamically irrelevant variables in the space of states; the elimination of such variables results in an interestingly distinct theory (the new phase space will be of smaller dimension than the original, and (unless one restricts attention to a subspace) will be a Poisson space even when the original phase space is symplectic).[36]

are $n!$ states upstairs for every state downstairs; but if we look at states in which exactly two particles share the same position and momentum, this factor goes down to $\frac{n!}{2}$.

35 Hence while there may perhaps be more temptation to apply the quotienting procedure across the board in the case of spaces of solutions than in the case of individual solutions, familiar haecceitistic modal intuitions – see Adams (1979) and Lewis (1986, section 4.4) – provide a countervailing force.

36 Gauge theories provide an especially vivid example. One starts with a theory in many ways analogous to that of M_0 above, then quotients out the action of the infinite dimensional group of gauge transformations. For details, see Belot (2003), and Earman (this volume, Part I), and Redhead (this volume).

Acknowledgements

I would like to thank the editors of this volume for their patience, encouragement, and helpful advice.

References

Adams, R. (1979). 'Primitive thisness and primitive identity'. *Journal of Philosophy*, **76**, 5–26.

Ashtekar, A., and Schilling, T. (1999). 'Geometrical formulations of quantum mechanics'. In *On Einstein's Path*, ed. A. Harvey. Berlin: Springer-Verlag.

Barbour, J., and Bertotti, B. (1982). 'Mach's principle and the structure of dynamical theories'. *Proceedings of the Royal Society of London A*, **382**, 295–306.

Belot, G. (2001). 'The principle of sufficient reason'. *Journal of Philosophy*, **98**, 55–74.

(2003). 'Symmetry and gauge freedom', *Studies in History and Philosophy of Modern Physics*, **34**, 189–225.

Brown, H., and Sypel, R. (1995). 'On the meaning of the relativity principle and other symmetries'. *International Studies in the Philosophy of Science*, **9**, 235–53.

Cornford, F. (1997). *Plato's Cosmology*. Indianapolis, IN: Hackett.

Deligne, P., and Freed, D. (1999). 'Classical field theory'. In *Quantum Fields and Strings: a Course for Mathematicians*, Vol. 1, ed. P. Deligne *et al*. Providence, RI: American Mathematical Society.

Descartes, R. (1991). *Principles of Philosophy*. Dordrecht: Kluwer.

Eddington, A. (1920). *Space, Time, and Gravitation*. Cambridge: Cambridge University Press.

Giulini, D. (2001). 'Uniqueness of simultaneity'. *British Journal for the Philosophy of Science*, **52**, 651–70.

Gödel, K. (1949a). 'An example of a new type of cosmological solution of Einstein's field equations of gravitation'. *Review of Modern Physics*, **21**, 447–50.

(1949b). 'A remark about the relationship between relativity theory and idealistic philosophy'. In *Albert Einstein: Philosopher–Scientist*, ed. P. Schilpp. New York: Harper and Row.

(1952). 'Rotating universes in general relativity theory'. *Proceedings of the International Congress of Mathematicians, Cambridge, MA, 1950*. Providence, RI: American Mathematical Society.

Hacking, I. (1975). 'The identity of indiscernibles'. *Journal of Philosophy*, **72**, 249–56.

(1978). 'On the reality of existence and identity'. *Canadian Journal of Philosophy*, **8**, 613–32.

(1979). 'What is logic?' *Journal of Philosophy*, **76**, 285–319.

Hodges, W. (1997). *A Shorter Model Theory*. Cambridge: Cambridge University Press.

Huggett, N. (1999). 'Atomic metaphysics'. *Journal of Philosophy*, **96**, 5–24.

Huygens, C. (1977). 'The motion of colliding bodies', translated by R. J. Blackwell. *Isis*, **68**, 574–97.

Ismael, J. (1997). 'Curie's principle', *Synthese*, **110**, 167–90.

Jeans, J. (1936). 'Man and the universe'. In *Scientific Progress*, ed. E. Appleton *et al*. New York: MacMillan.

Kramer, D., Stephani, H., Herlt, E., and MacCallum, M. (1980). *Exact Solutions of Einstein's Field Equations*. Cambridge: Cambridge University Press.

Landsman, N. (1998). *Mathematical Topics Between Classical and Quantum Mechanics*. Berlin: Springer-Verlag.

Lewis, D. (1986). *On the Plurality of Worlds*. Oxford: Blackwell.

Malament, D. (1977). 'Causal theories of time and the conventionality of simultaneity'. *Noûs*, **11**, 293–300.

Marsden, J. (1992). *Lectures on Mechanics*. Cambridge: Cambridge University Press.

Marsden, J., and Ratiu, T. (1994). *Introduction to Mechanics and Symmetry*. Berlin: Springer-Verlag.

Minkowski, H. (1952). 'Space and time'. In *The Principle of Relativity*, ed. A. Einstein *et al.* New York: Dover.

Olver, P. (1993). *Applications of Lie Groups to Differential Equations*. Berlin: Springer-Verlag.

Singer, S. (2001). *Symmetry and Mechanics*. Basel: Birkhäuser.

Stein, H. (1991). 'On relativity and openness of the future'. *Philosophy of Science*, **58**, 147–67.

ter Haar, D. (1995). *Elements of Statistical Mechanics*, 3rd edn. Oxford: Butterworth-Heinemann.