The Metaphysics of Invariance

David Schroeren*

April 18, 2018

Abstract

Fundamental physics contains an important link between properties of elementary particles and continuous symmetries of particle systems. For example, properties such as mass and spin are said to be 'associated' with specific continuous symmetries. These 'associations' have played a key role in the discovery of various new particle kinds, but more importantly: they are thought to provide a deep insight into the nature of physical reality.

The link between properties and symmetries has been said to call for a radical revision of perceived metaphysical orthodoxy. However, if we are to use claims about an 'association' between properties and symmetries in the articulation of metaphysical views, we first need to develop a sufficiently precise understanding of the content of these claims. The goal of this paper is to do just that.

1 Introduction

Fundamental physics contains an important link between properties of elementary particles and continuous symmetries of particle systems. As Steven Weinberg puts it, properties such as mass and spin are each "associated with" specific continuous symmetries.[1] This link has played a key role in the discovery of various new particle kinds, but more importantly: it is thought to provide a deep insight into the nature of physical reality. For example, Weinberg claims that the relevant

*The Department of Philosophy, 1879 Hall, Princeton University, Princeton, NJ 08544-1006. E-mail: schroeren@princeton.edu.

[1]Weinberg (1993, p. 146) Other physicists have made similar claims; for example, Heisenberg (1976, p. 924), Penrose (2004, p. 568), and Ne'eman and Sternberg (1991, p. 327) The link between properties and symmetries was first hypothesised by Eugene Wigner (1939).

physical properties "are what they are *because* of the [associated] symmetries of the laws of nature."[2]

According to an influential line of scholarship, the link between properties and symmetries calls for a radical revision of metaphysics—away from a perceived orthodoxy according to which the world fundamentally consists of objects instantiating properties, and toward a metaphysics which conceives of fundamental reality as purely 'structural' in a sense to be elucidated (at least in part) by appeal to just this link between properties and symmetries. Proponents of this line—also known as *ontic structural realism*—claim that the 'association' between physical properties and symmetries is "the basis of the claim that such properties should be conceived of structurally,"[3] in the sense that "the properties that particles have [...] appear to be explicable via considerations of [symmetry] structure" so that "what should be regarded as properly fundamental is the symmetry structure [...] that explains [the properties of elementary particles]."[4] However, if we are to use claims about the 'association' between physical properties and continuous symmetries in the articulation of metaphysical theories, let alone as the basis for a radical revision of metaphysics, we first need to develop a sufficiently precise understanding of the content of these claims. The purpose of this paper is to do just that.

It may be tempting to think that the content of the 'association' between physical properties and symmetries is the more familiar claim, known as Noether's theorem, that every continuous symmetry corresponds to a *conserved quantity*.[5,6]

---

[2](Weinberg, 1993, p. 138n, emphasis added). Elsewhere, Weinberg claims that the 'association' between properties and symmetries shows that "at the deepest level, all we find are symmetries and responses to symmetries" (Weinberg, 1987, p. 80) and that "matter thus loses its central role in physics" (Weinberg, 1993, p. 139).

[3](French, 2014, p. 109).

[4](McKenzie, 2014, p. 21). This is echoed by other authors; e.g. Ladyman (1998); Ladyman and Ross (2007); Kantorovich (2003). Ontic structuralists also use the link between properties and symmetries to argue *against* other metaphysical views such as certain varieties of dispositional essentialism; for example, French (2014, pp. 251) and Saatsi (2017); cf. (Livanios, 2010).

[5](Noether, 1918). Nothing in this paper hinges on the specifics of how the notion of 'symmetry' is understood. For definiteness, think of a symmetry of a physical system as a transformation of a given type (such as rotations) that *at least* preserves the laws of the system; for an overview of various accounts of symmetry, see (Dasgupta, 2016).

[6]Physical quantities (like mass and energy) can be thought of as associated with a range of *values*, each of which is a property in its own right.

For example, Noether's theorem entails that rigid rotations about a given point are a symmetry of a physical system iff the system's total angular momentum about that point is conserved, i.e. constant in time—a fact which explains why figure skaters spin more quickly when drawing in their arms. But this is clearly not the claim physicists have in mind when they say that a given property is 'associated' with a continuous symmetry. Whereas this 'association' concerns "the way particles behave when you perform various *symmetry transformations*" such as rotations,[7] Noether's theorem concerns the way particles behave *under temporal evolution*, i.e. whether certain physical quantities are conserved. Noether's theorem is therefore not what is meant by the claim that a property is 'associated' with a continuous symmetry.

Another tempting thought is that a physical property is 'associated' with a continuous symmetry iff this property is *invariant* under that symmetry, in the sense that any two states of the system related by that symmetry agree about this property. There is something right about this proposal: every property 'associated' with a symmetry is invariant under that symmetry. But there is also something wrong about this proposal: a property can be invariant under a symmetry without being 'associated' with that symmetry. Consider a classical n-particle system in three-dimensional Euclidean space. At every instant of time, the system has values of various *vector* quantities, such as total linear momentum $\vec{p}$ and, for every spatial point, the total angular momentum $\vec{L}$ of the system about that point. The system also has values of various *scalar* quantities, such as the magnitude of total linear momentum $|\vec{p}|$ and, for any given point, the magnitude $|\vec{L}|$ of total angular momentum about that point. Both of these magnitudes are invariant under rigid rotations about any given point.[8] And yet, for every point x, only the magnitude of total angular momentum about x is 'associated' with rigid rotations about x in the sense identified by physicists.

The question we're asking has a standard mathematical answer—one that physicists like Weinberg presumably have in mind when making the more colourful

[7](Weinberg, 1987, p. 79); emphasis added.

[8]Throughout this paper, I understand the relevant transformations *actively*; that is, as transformations on physical systems, rather than merely as transformations of the mathematical devices representing those systems. See (Brading and Castellani, 2007, pp. 1342-3) for more on this distinction.

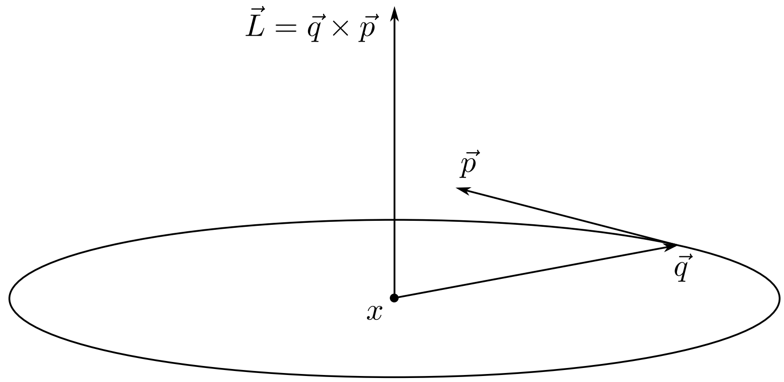

Figure 1: Illustration of a classical one-particle system in orbital motion about x. $\vec{q}$ represents the position of the particle in a co-ordinate system originating at x, $\vec{p}$ its linear momentum and $\vec{L} = \vec{q} \times \vec{p}$ its angular momentum about x.

claims cited above. This answer goes roughly as follows. The notion of a continuous symmetry is mathematically formalised by the notion of a *group*: roughly, a collection of transformations of a given type, such as rotations. And every group G of continuous symmetries of a system is associated with a mathematical space capturing certain essential features of the transformations in G.[9] Now, there are mathematical objects (functions in the context of classical mechanics, operators in the context of quantum mechanics) whose invariance under G can be derived exclusively from the essential features of transformations in G captured by this space—objects which are called *Casimir invariants* of G and whose number depends on G. The key observation is now that each property 'associated' with a particular symmetry group is represented by one of the Casimir invariants of that group.[10] For example, the magnitude of total angular momentum about spatial point x of a classical n-particle system is represented by a Casimir invariant of the group of rigid rotations of the system about x, whereas the magnitude of total linear momentum is not. Similarly, mass and spin of elementary particles are each represented by one of the two Casimir invariants of the group of relativistic boosts, rotations, and

[9]This space is called the *Lie algebra* of G.

[10]By 'representation' I mean the way in which items of a linguistic or mathematical sort (theories, models, sentences, variables, predicates, functions, operators) are used to make claims about items of a 'metaphysical' or 'worldly' sort (worlds, facts, propositions, objects, properties, quantities).

translations.[11] The standard mathematical answer is thus that a physical property P is 'associated' with a continuous symmetry group G iff the mathematical object representing P is a Casimir invariant of G.

A mathematical physicist or a certain sort of philosopher might think that this is all that needs to be said about the matter. According to this viewpoint, mathematical physics in itself is perfectly adequate and fully intelligible as a representation of the underlying reality, and no further understanding is to be gained by going 'beyond' physics into the murky waters of metaphysics.[12]

But for another kind of philosopher, the standard mathematical answer isn't good enough. According to this philosopher, there is something dissatisfying about characterising features of reality by reference to the mathematical properties of their mathematical representations. If someone unfamiliar with the concept of money asks what is conveyed about a bag by saying that it contains some amount of US dollars in cash, the response 'that the bag has whatever feature is accurately represented by the dollar sign printed on its side' is likely to be dismissed as unilluminating. It is similarly inadequate to characterise the feature that a physical property has just in case it is 'associated' with a symmetry group as whatever feature of this property is accurately represented by the fact that its mathematical representative is a Casimir invariant of that group. Therefore, once the mathematical physicist has explained the notion of a Casimir invariant, there's a further question that this philosopher would like to see addressed: what information *about a property* is conveyed by saying that its mathematical representative—the function or operator associated with it—is a Casimir invariant of some continuous symmetry group? This is the question I am going to answer in this paper.

Even if it is taken for granted that the standard mathematical answer is insufficient, one might still wonder how far beyond the mathematical physics one needs to go for an account of this link to be illuminating. I don't have an *a priori* standard relative to which I can answer this question, and relative to which the standard mathematical account is insufficient. All I can do is to present my account

[11] This group is referred to as the *Poincaré group*.

[12] One influential philosopher of physics who holds this view is David Wallace: "I'm happy to say that the way theories represent the world is inherently mathematical, so that the reason property P is 'associated' with symmetry G is simply because P is represented by the Casimir invariant of G, and there's no need to say more." (Personal correspondence, quoted with permission.)

of the ‘association’ between physical properties and symmetries—an account that will make clear exactly what I take to be missing from the standard mathematical answer.

One consequence of the approach I take in this paper is that I distinguish sharply between features of reality and the mathematical objects representing those features. This means that there will usually not be a straightforward route from a feature of the mathematical representative of a property to a feature of this property.

I proceed as follows. I explain the standard mathematical answer in section (2) by considering the case study of total angular momentum about a point x of a classical n-particle system and the group of rotations of the system about x. Subsequently, I explain the metaphysical content of the notion of a Casimir invariant. According to the proposal I develop in section (3), what is conveyed about a property by the fact that it is represented by a Casimir invariant of a given group is that this property is a *fine-grained invariant* under that group—a notion I introduce. This notion is ‘fine-grained’ in the sense that not every invariant (in the ordinary sense) under a group is also a fine-grained invariant under that group: for example, it will turn out that the magnitude of total angular momentum about x is a fine-grained invariant under rigid rotations about x, whereas the magnitude of total linear momentum is not.

2 Casimir Invariance

According to the standard mathematical answer, a physical property P is ‘associated’ with a continuous symmetry group G iff the mathematical object representing P is a *Casimir invariant* of G. The goal of this section is to introduce this notion.

2.1 The Framework

I do so by considering a classical n-particle system described in terms of *Hamiltonian mechanics*. Here, the goal is to explain what it is for the mathematical representative of the magnitude of total angular momentum about spatial point x to be a Casimir invariant of the group of rotations of the system about x.

Why Hamiltonian mechanics? One important reason is that the Hamiltonian treatment of continuous symmetries is most similar to the corresponding formalism in quantum theory, including quantum field theory. Hamiltonian mechanics is therefore a natural setting in which to explain what it is for a property to be 'associated' with a symmetry group in the way that physicists have identified in the context of quantum theory, but without being distracted by the technical difficulties of quantum theory.[13]

First, let me introduce a few basic features of Hamiltonian mechanics. In this framework, the state space of a system—the space of *possible instantaneous states* of the system—is represented by a mathematical space called *phase space*.[14] For example, consider a system of n particles in three-dimensional Euclidean space represented by $\mathbb{R}^3$. The phase space of this system is given by its configuration space—the space whose points each represent a possible instantaneous spatial configuration of the system—together with an n-dimensional vector space attached to every point, a space whose elements each represent a possible assignment of a specific value of linear momentum to every particle at that point.[15] We can describe any given state of the system by a $6n$-tuple of numbers with regard to some co-ordinate system of $\mathbb{R}^3$: three spatial co-ordinates as well as three momentum co-ordinates for each of the n particles. In this way, Hamiltonian mechanics characterises each state of a system by two sorts of basic, independent features: the positions and momenta of each constituent particle.

Each physical quantity of the system is associated with a smooth distribution of values over state space; that is, with a smooth assignment, to each possible state, of a value of this quantity. These distributions are specified by smooth functions from phase space to the real numbers (in the case of scalar quantities) or by smooth functions from phase space to triples of real numbers (in the case

[13] To be sure: there are *other* respects in which quantum theory has a greater resemblance with Lagrangian mechanics than with Hamiltonian mechanics—particularly with regard to path integral methods (Peskin and Schroeder, 1995, Chpt. 9).

[14] What are 'possible instantaneous states'? That depends on one's preferred view about the philosophy of time. On an A-theoretic picture of time, we can think of possible instantaneous states as possible worlds. On a B-theoretic picture of time, we can think of possible instantaneous states as maximal properties of possible worlds at times.

[15] In mathematical terms, the phase space of an n-particle system is the *cotangent bundle* of its configuration space.

of vector quantities). For example, the distribution of total energy over the state space of the n-particle system is represented by a real-valued function referred to as *the* Hamiltonian and denoted by H.[16] For the sake of brevity, I will abbreviate 'the state space distribution associated with a quantity Q is represented by the phase space function f_Q' by 'Q is represented by f_Q'.[17]

I will take for granted that physical quantities can be distinct despite being represented by the same phase space function. Call two quantities Q, Q' *state-space co-extensive* just in case their associated distributions assign the same value to every state in state space, i.e. just in case the corresponding phase space functions f_Q, $f_{Q'}$ are mathematically co-extensive: $f_Q = f_{Q'}$. The claim is then that there are distinct state-space co-extensive quantities.

One motivation for this assumption is that state-space co-extensive quantities need not be *metaphysically* co-extensive. More precisely: suppose we think of the distributions associated with physical quantities as assigning values to every *metaphysically possible state* of the system, rather than merely to the possible states represented by points in phase space. Then, for any phase space function, there will generally be distinct quantities such that the restriction to state space of their associated distributions is accurately represented by this function.

By way of example, suppose that total energy of the n-particle system is in fact state-space co-extensive with the quantity whose value at a state is the result of a certain specific operation on the position and momentum facts at that state corresponding to the sum of the values of kinetic energy and Newtonian gravitational energy at that state—a fact we may express in terms of the equation $H = T + V_g$, where T, and V_g are phase space functions representing kinetic and

[16]Physical quantities associated with smooth phase space functions are generally not invariant under changes of inertial reference frames (or *Galilei frames*): for example, there are Galilei frames which disagree about linear momentum and total energy. This means that we should always be talking about these quantities *relative to a particular Galilei frame*. However, the way these quantities are talked about in physics and philosophy often doesn't make their Galilei-dependence explicit. In this paper, I follow this practice for the sake of brevity.

[17]Not all quantities represented by smooth phase space functions are equally natural: any 'gerrymandered' function of position and momentum, such as the product $\vec{q}^{\,i} \cdot \vec{p}^{\,k}$ of the position vector of the i-th particle with the linear momentum vector of the k-th particle, is less natural than (for example) the function which represents the linear momentum of a given particle. This point is familiar from the metaphysics of properties: not every set of individuals is equally natural; and only few, privileged sets are perfectly natural; cf. (Lewis, 1986, pp. 59).

Newtonian gravitational energy, respectively. But total energy and the quantity represented by $T+V_g$ aren't *metaphysically* co-extensive: for example, they disagree at metaphysically possible states which contain *electrostatic energy*. Thus, there are distinct state-space co-extensive quantities.

This paper focuses on the quantity of *total angular momentum* about a given spatial point x. The state space distribution of this quantity is specified by the vector function $\vec{L} = \vec{L}^1 + ... + \vec{L}^n$, where $\vec{L}^k = \vec{q}^{\,k} \times \vec{p}^{\,k}$ represents the angular momentum of the k-th particle about x, $\vec{q}^{\,k}$ its position in a co-ordinate system centered on x, and $\vec{p}^{\,k}$ its linear momentum. The scalar components of $\vec{L}$ are real-valued phase space functions each of which represents the state space distribution of the component of angular momentum along *some* axes through x.[18]

Certain continuous sequences of states (or *state space curves*) capture important dynamical and modal facts about the system. Most importantly, the state space curves that count as *possible dynamical histories* of the system allow us to answer questions such as 'what *will* the system be like in five minutes'? Another important type of state space curves allow us to answer questions such as 'what *would* the system be like if it *were* rigidly rotated by 180° about an axis through some spatial point x?' These state space curves correspond to *rigid rotations* of the system about axes intersecting x.

State space curves of a given type are specified in terms of some particular *flow* on phase space, i.e. in terms of some particular family of phase space curves such that each phase space point is intersected by exactly one curve in this family. (Think of a flow as the family of curves traced out by water molecules of a river.) Each flow has a real parameter: for example, for every real number t, the *dynamical* flow representing the dynamical histories of the system assigns to each phase space point its image under dynamical evolution by t units of time. Similarly, for each real number ϕ, the *rotational* flow corresponding to rigid rotations of the system about an axis through x assigns to each phase space point its image under a rigid rotation by an angle of ϕ about this axis.

It is worth stressing that the relevant rotations are *non-dynamical*: whereas the parameter of the dynamical flow is interpreted as time, the parameter of a

[18]Let $\{e_1, e_2, ..., e_n\}$ be an orthonormal basis of a real vector space V. The *scalar components* of any $v = v_1e_1 + v_2e_2 + ... + v_ne_n \in V$ are the real numbers v_1, v_2,..., v_n.

given rotational flow is interpreted as the corresponding angle of rotation; and states related by a rotation are generally not related by dynamical evolution, as illustrated in figure (2).

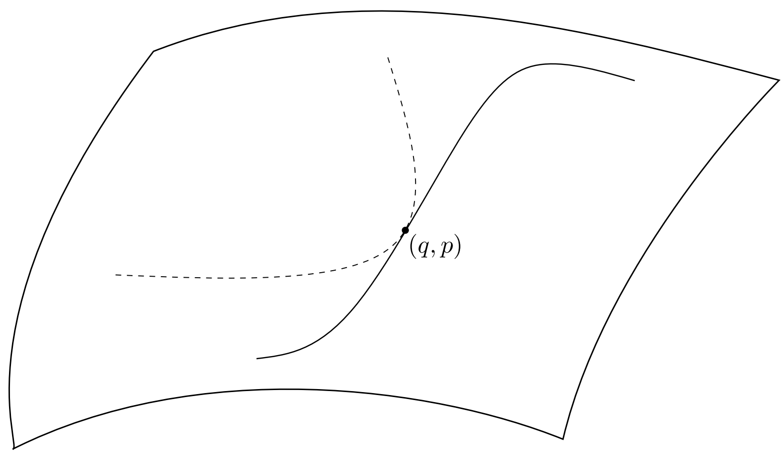

Figure 2: Illustration of phase space with a dynamical flow line and a rotational flow line (dotted) through phase space point (q, p).

This is where a crucial feature of Hamiltonian mechanics comes into play: each of these flows is *determined* by a smooth function on phase space in a way that *preserves* this function. A flow which is determined by and which preserves a smooth function is referred to as the *Hamiltonian flow* of this function; and every smooth function is referred to as a *generator* of its Hamiltonian flow. For example, the dynamical flow is the Hamiltonian flow generated by the total energy function H, and the rotational flow corresponding to rigid rotations about some axis through x is given by the scalar component of $\vec{L}$ along this axis.

The notion of a generator of a Hamiltonian flow is key to this paper, so bear with me while I explain it.

2.2 Hamiltonian Flows and their Generators

What is it for a smooth function to be a generator of its Hamiltonian flow? Here is an informal picture. We are familiar with the sense in which a *vector field* generates a flow. (Think of a vector field on some space as a function which assigns a vector to every point in this space.) For example, the flow of the velocity vector field of a river can be thought of as the family of curves traced out by the water molecules dragged along by the current of the river. But what is it for a *scalar* function to be a generator of a flow?

There's a fairly intuitive sense in which a scalar function is the generator of the flow of its *gradient*. Recall: the gradient of a smooth function f on a Riemannian manifold is the vector field which assigns to each point of the manifold the vector that points in the direction of greatest increase of f at that point.[19] And every smooth function f on such a manifold determines its gradient in a two-step process. The first step consists in an operation that takes smooth functions as input and throws away all but two specific sorts of information: first, information about the regions of the manifold in which the input function is constant, also called *level surfaces* of this function;[20] and second, information about the differences in value of this function between each of its level surfaces. In mathematics jargon, this operation is called the *exterior derivative*. For example, the information contained in the exterior derivative of the elevation function of a piece of terrain suffices to draw an accurate map of the terrain: on the basis of this information, one can draw contour lines as well as the differences in elevation between these lines. What would be missing from the map, however—the information thrown away by the exterior derivative—is the absolute elevation of a given point of the terrain.

The information contained in the exterior derivative of f is not yet packaged in the form of a vector field. This is the task of the second step, in which the exterior derivative of f is turned into the *gradient* of f. The specifics of this mathematical procedure are not relevant for our purposes, except for the fact that it is defined with regard to the Riemannian metric on the underlying manifold, a rule we can think of as the Euclidean 'dot product' of two vectors.

It is by way of these two steps that a function determines its gradient. But as noted above, a vector field also determines its flow. Thus, f also determines the flow of its gradient, which is captured by saying that f is a *generator* of this flow.[21]

But what is it for a function on phase space to generate its *Hamiltonian* flow?

[19] A Riemannian manifold is a smooth manifold equipped with a *Riemannian metric*, a positive-definite map from pairs of vector fields to smooth functions that we can think of as a generalisation of the Euclidean dot product.

[20] Level surfaces are generalisations of the one-dimensional contour lines familiar from hiking maps, i.e. the lines of constant elevation of the relevant piece of terrain.

[21] Since there are distinct functions with the same exterior derivative, the flow of each gradient vector field has several generators.

The answer lies in a different way to package the information contained in the exterior derivative, a different procedure for turning this information into a *vector field*—a procedure peculiar to Hamiltonian mechanics and defined with regard to a rule known as the *symplectic form* on phase space. The key difference between the procedure defined in terms of the Riemannian metric and the procedure defined in terms of the symplectic form is this: whereas the former turns the exterior derivative of a function into a vector field that everywhere points in the direction of greatest increase of the function, the latter turns the exterior derivative into a vector field that everywhere points *along the level surfaces* of this function, as illustrated in figure (3). Vector fields of this sort are referred to as *Hamiltonian*

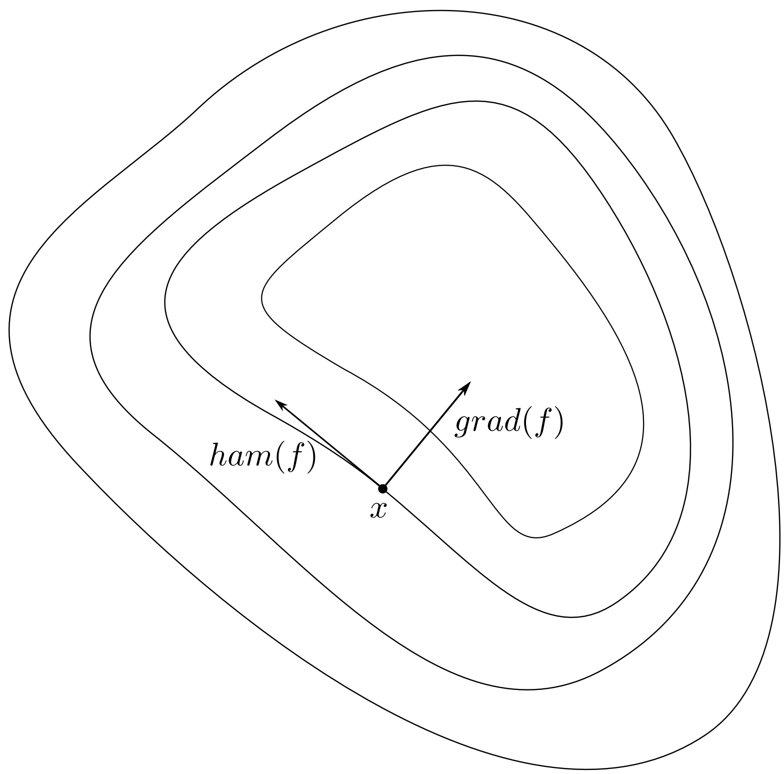

Figure 3: Illustration of the level surfaces (contour lines) of a function f on a two-dimensional manifold. $grad(f)$ and $ham(f)$ denote the values at x of the gradient and the Hamiltonian vector field of f, respectively.

.

vector fields. As a result, every flow line of the Hamiltonian vector field of f is confined to a particular level surface of f, and so f is invariant along this flow.[22] We thus understand what it is for a scalar function to generate its Hamiltonian flow, and thus what it is for the component of $\vec{L}$ along some axis through spatial point x to be a generator of the Hamiltonian flow corresponding to rigid rotations about this axis.

Now, the rigid rotations about arbitrary axes intersecting x *taken together* have the structure of a group, namely the group $SO(3)$ of orientation-preserving isometric rotations in three-dimensional Euclidean space.[23] More precisely, rigid rotations of the n-particle system about a given point are implemented by an *action* of $SO(3)$ on the phase space of the system, i.e. by a rule for how the position and momentum of each constituent particle transforms under every abstract element of the group: for every spatial point, the action of $SO(3)$ as rotations about that point consists in an assignment, to each abstract element of $SO(3)$, of a rotation by *some* angle about *some* axis intersecting that point.

This has an important consequence: any two phase space points related by a rotation about spatial point x are related by a rotation about *some specific axis* i intersecting x. *A fortiori*, any two such phase space points lie on a flow line belonging to the Hamiltonian flow generated by the i-component of $\vec{L}$. In this way, the action of $SO(3)$ as rigid rotations about x is implemented by Hamiltonian flows corresponding to rotations about axes through x, and so the generators of the flows about these axes—the components of $\vec{L}$ along these axes—are also generators of this action of $SO(3)$, or *SO(3)-generators*, for short.

Although each $SO(3)$-generator is invariant under its own Hamiltonian flow, $SO(3)$-generators are not $SO(3)$-invariant: it is not the case that any two phase space points related by a rotation about some axis through x agree about every component of $\vec{L}$. Exactly how do the components of $\vec{L}$ change under rotations? This is the last question we need to answer before we can state the notion of a

[22] Since there are distinct functions with the same exterior derivative, each Hamiltonian flow has several generators.

[23] The statement that rotations about a point have the structure of a *group* is equivalent to the conjunction of four claims. First, the combination of any two rotations is also a rotation. Second, the result of first performing rotation R followed by the combination of rotations R' and R'' is the same as first performing the combination of R and R', followed by the rotation R''. Third, every rotation can be reversed. Fourth, some rotations are equivalent to doing nothing—e.g. rotations by 0 and 360 degrees.

Casimir invariant.

2.3 $SO(3)$ Poisson brackets

Physics is often concerned with determining the *rates of change* of physical quantities under certain important variations of the system. For example, *velocity* is the instantaneous rate of change through time of position, and *net force* is the instantaneous rate of change through time of total linear momentum.

In Hamiltonian mechanics, facts of this sort are represented by the derivatives of phase space functions with regard to the parameter of the Hamiltonian flow that characterises the relevant variation of the system. The instantaneous rate of change *under dynamical evolution* of the position of the k-th particle is given by the derivative of $\vec{q}^{\,k}$ with regard to *time*, the parameter of the Hamiltonian flow generated by the total energy function H. Similarly, the rate of change of the linear momentum of the k-th particle *under rotations* about some axis through x is represented by the derivative of $\vec{p}^{\,k}$ with *with regard to the angle of rotation* about that axis, the parameter of the Hamiltonian flow generated by the component of $\vec{L}$ along that axis.

The derivative of a smooth function with regard to the parameter of a Hamiltonian flow can be expressed directly in terms of the generator of this flow. More specifically, the derivative of a function f with regard to the parameter λ of the Hamiltonian flow generated by another function g is given by[24]

$$\frac{df}{d\lambda} = \{f, g\}, \tag{1}$$

where $\{\cdot\,,\cdot\}$ is the *Poisson bracket*, a product on the space of smooth functions satisfying certain conditions, the details of which need not concern us here.[25]

For present purposes, the following intuitive picture of the mathematical content of (1) shall suffice. A perfectly standard way to represent the rate of change of the elevation of a piece of terrain along a path through the terrain is by means of the

[24] Recall: we're supposing that the f and g are real-valued functions on phase space *only*. This entails that $\frac{\partial f}{\partial \lambda} = 0$, where λ is the parameter of the flow of g.

[25] The Poisson bracket is a *Lie bracket*; that is, it is anti-symmetric, bi-linear, non-associative and satisfies the Jacobi identity: for any smooth functions f, g, h, $\{\{f, g\}, h\} + \{\{g, h\}, f\} + \{\{h, f\}, g\} = 0$.

derivative of the elevation function in the direction of this path.[26]

Suppose now that we are considering an entire family of paths through the terrain, and suppose that this family can be represented by the flow of some vector field Y. We then have a notion of a derivative of the elevation function in the direction of Y. Now, this derivative can be thought of as the function which assigns to every point of the terrain a number that tells us how much of the greatest increase of elevation we're experiencing in the direction of Y—the function whose value at a point x is the *projection*, with regard to the Riemannian metric, of the value of the gradient of the elevation function at x onto the value of Y at x, as illustrated in figure (4).

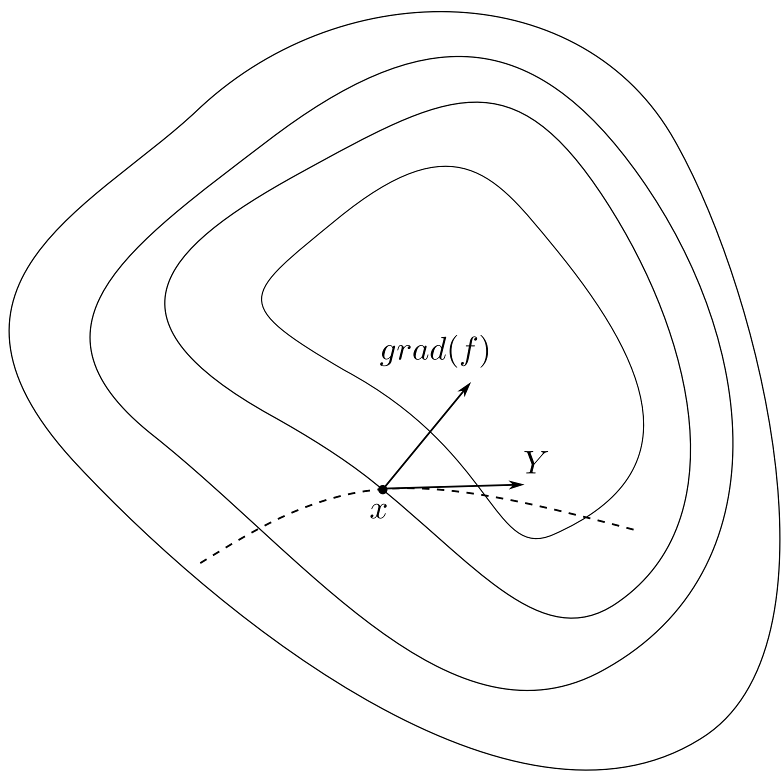

Figure 4: Illustration of a function f via its contour lines, showing a flow line (dotted) of the vector field Y and the value $grad(f)$ of the gradient of f at x.

Equation (1) has a very similar mathematical role: $\{f, g\}$ also captures a sort of derivative of f in the direction of a vector field—namely, in the direction of the

[26]In mathematical jargon, this derivative is known as the *directional derivative* of the elevation function along the path.

Hamiltonian vector field of g. However, the vector field that is being projected onto the Hamiltonian vector field of g is not the gradient of f, but rather the Hamiltonian vector field of f. Moreover, the projection operation is different: whereas in the previous case the projection operation is given by the Riemannian metric, here it is given by the symplectic form. Notwithstanding these differences, we may think of the Poisson bracket $\{f, g\}$ as a derivative of f *in the direction of the Hamiltonian vector field of* g, and thus as a phase space function that captures how f *varies* along the Hamiltonian flow of g.

Poisson brackets are therefore the tool we need to determine how the components of $\vec{L}$ vary under rotations about some spatial point x. Denote by L_i the scalar component of $\vec{L}$ along axis i through x, and let 1, 2, and 3 be the labels of pairwise perpendicular axes through x. Then, $\{L_1, L_2\}$ is the derivative of L_1 along the Hamiltonian flow generated by L_2. And as it turns out, the derivative of any given $SO(3)$-generator along the flow of another $SO(3)$-generator is also an $SO(3)$-generator: $\{L_1, L_2\} = L_3$. Equations of this sort are called *SO(3)-Poisson bracket equations* and summarised as

$$\{L_i, L_j\} = \epsilon_{ijk} L_k, \tag{2}$$

where i, j, k take values in the labels of any given triple of pairwise perpendicular axes intersecting x and ϵ_{ijk} is 1 for even permutations of ijk, -1 for uneven permutations and 0 otherwise.[27]

We need to take note of a related, important use of Poisson brackets. It might be remembered from high school calculus that a function of a single variable is constant just in case its derivative vanishes everywhere. Similarly, a function f is constant (or invariant) along the Hamiltonian flow of another function g iff the derivative of f along this flow vanishes everywhere in phase space; that is, iff $\{f, g\} = 0$. Facts about the invariance of functions along Hamiltonian flows are thus facts about vanishing Poisson brackets.

We can use this to say what it is for a function to be invariant under rotations about some spatial point x: a function f has this property just in case, for every

[27] In mathematical jargon, equation (2) means that the set of $SO(3)$-generators forms a *Lie algebra* with regard to the Poisson bracket, called the *SO(3)-Poisson Lie algebra*.

axis i through x, its derivative along the Hamiltonian flow corresponding to rigid rotations about i vanishes: $\{f, L_i\} = 0$ for every axis i through x. I refer to functions which have this property as *SO(3)-constant* functions.

Note that the physical quantities represented by $SO(3)$-constants are such that any two states related by an $SO(3)$-transformation agree about these quantities—which is to say that these quantities are $SO(3)$-invariant. Whenever there is a risk of confusion, I refer to the quantities represented by $SO(3)$-constant functions as *ordinary* $SO(3)$-invariants.

The goal of this section was to explain how the $SO(3)$-generators change under rotations. As promised, this was the last thing we needed to understand before articulating the notion of a Casimir invariant. To this I now turn.

2.4 Casimir Invariants

I follow the standard definition in physics texts.[28] For a smooth function C on phase space to be a *Casimir invariant* of the action of $SO(3)$ as rigid rotations about x (or an $SO(3)$-*Casimir*, for short) is for C to satisfy two conditions. First, C must be an $SO(3)$-constant: $\{C, L_i\} = 0$ for every axis i through x. Second, C must be a function exclusively of the $SO(3)$-generators.

One phase space function which satisfies these conditions is $L^2 = L_1^2 + L_2^2 + L_3^2$, the Euclidean dot product of the angular momentum vector function with itself. First, L^2 is an $SO(3)$-constant; that is, $\{L^2, L_i\} = 0$ for every axis i through x.[29] Second, L^2 is a function only of $SO(3)$-generators, as can be seen by noting that no two states can differ about L^2 without differing about at least one component of $\vec{L}$. Importantly, L^2 is the unique 'independent' $SO(3)$-Casimir, in the sense that every other $SO(3)$-Casimir is a function of L^2—for example, the function $cos(L^2)$. For this reason, we may refer to L^2 as *the* $SO(3)$-Casimir.

L^2 has a systematic relationship with the function $|\vec{L}|$ representing the magnitude of angular momentum M: $|\vec{L}|$ is the square root of L^2. In the remainder of this paper, I follow an informal convention in physics by treating L^2 and $|\vec{L}|$ as having the same representational content, so that M may be said to be represented by L^2 as much as by $|\vec{L}|$. In so doing, it is important to keep in mind that any given

[28](Penrose, 2004, pp. 553,568), (Fuchs and Schweigert, 2003, pp. 254).

[29]This is derived in (Goldstein et al., 2001, p. 418).

value of M is represented by the square root of the corresponding numerical value of L^2.

The functional dependence of $SO(3)$-Casimirs on $SO(3)$-generators is responsible for an important feature of Casimirs: a function C is an $SO(3)$-Casimir iff the fact that C is $SO(3)$-constant can be derived from the $SO(3)$-Poisson brackets alone. In particular, the $SO(3)$-Poisson brackets $\{L_i, L_j\} = \epsilon_{ijk}L_k$ are sufficient to derive the equations $\{L^2, L_i\} = 0$. This feature of $SO(3)$-Casimirs will be important later, so keep it in mind.

Not every $SO(3)$-constant is also an $SO(3)$-Casimir. For example, the function p^2 representing the magnitude of total linear momentum is an $SO(3)$-constant: it obeys $\{p^2, L_i\} = 0$ for every relevant axis i and thus satisfies the first condition of being an $SO(3)$-Casimir. But p^2 is not a function of the $SO(3)$-generators: the phase space points $(\vec{q}, \vec{p})$ and $(2\vec{q}, \frac{1}{2}\vec{p})$ representing states of a one-particle system agree about $\vec{L} = \vec{q} \times \vec{p}$ and thus about all components thereof, but disagree about p^2. It follows that p^2 is not a function of the $SO(3)$-generators, and thus not an $SO(3)$-Casimir. This is confirmed by the fact that the $SO(3)$-constancy of p^2 cannot be derived from the $SO(3)$-Poisson brackets alone.

According to the standard mathematical account of the link between properties and symmetries, the magnitude of total angular momentum about x is 'associated' with the group $SO(3)$ of rigid rotations about x just in case the mathematical representative of this quantity is an $SO(3)$-Casimir. We now understand what this means: to be an $SO(3)$-Casimir is to be an $SO(3)$-constant function of the $SO(3)$-generators. The next task is to determine what is conveyed *about the magnitude of total angular momentum about* x by saying that its mathematical representative is an $SO(3)$-Casimir invariant.

3 Fine-Grained Invariance

Let me dispense with a tempting proposal immediately. As we just saw, one feature that distinguishes $SO(3)$-Casimirs from mere $SO(3)$-constants is that the former, but not the latter, are functions only of the $SO(3)$-generators. This is reflected in the fact that, for any triple of co-ordinate axes through x, L^2 can be written as the sum of squares of components of $\vec{L}$ along these axes, whereas the mere

$SO(3)$-constant p^2 cannot. It may now be tempting to think that it is a very similar feature that distinguishes the $SO(3)$-invariant *quantity* represented by the $SO(3)$-Casimir from other $SO(3)$-invariant quantities. Observe that the magnitude of total angular momentum M can be thought of as having the following sort of functional dependence on the components of angular momentum: M is the quantity such that, for any triple of pairwise perpendicular axes through x, the value of M at a state is the result of a specific operation on the position and momentum facts at that state corresponding to the square root of the sum of squares of the values of the angular momentum components along these axes at that state. According to the tempting proposal, what is conveyed about M by saying that its mathematical representative is an $SO(3)$-Casimir is the conjunction of two claims: first, M is an ordinary $SO(3)$-invariant; and second, M has just this sort of functional dependence on the components of angular momentum.

The problem with this proposal is that it cannot be applied in quantum mechanics. For example, actions of $SO(3)$ as rigid spatial rotations of a non-relativistic n-particle quantum system about a spatial point are generated by self-adjoint operators $\hat{L}_i$ that represent components of the total *orbital* angular momentum about that point. In this setting, the $SO(3)$-Poisson bracket equations (2) become the *commutator bracket* equations $[\hat{L}_i, \hat{L}_j] = \epsilon_{ijk}\hat{L}_k$, and the Casimir invariant associated with this action of $SO(3)$ is the operator that can be written as $\hat{L}^2 = \hat{L}_1^2 + \hat{L}_2^2 + \hat{L}_3^2$, interpreted as the magnitude of total orbital angular momentum.

The crux is now that it is impossible for a quantum system to simultaneously possess values of quantities represented by non-commuting operators.[30] This means that a quantum system cannot simultaneously have values for every component of total orbital angular momentum, and so no physical quantity can functionally depend on these components in the way fleshed out above.

The simple idea therefore does not carry over to quantum theory. But it is precisely the context of quantum theory in which the link between properties and symmetries is of greatest physical significance, and so any plausible account of this

[30]This follows from what is known as the *Kochen-Specker theorem* (Kochen and Specker, 1967) on the assumption that systems possess values of quantities independently of the context in which these quantities are measured.

link ought to be applicable to this context. Since the simple idea does not meet this requirement, it should be rejected.

What, then, is the right account? The simple idea we just considered hypostatises one mathematical feature of Casimir invariants: the functional dependence of Casimir invariants on the relevant generators. I propose that we focus on *another* distinctive feature of Casimir invariants: the fact that the $SO(3)$-constancy of $SO(3)$-Casimirs is derivable exclusively from the $SO(3)$-Poisson bracket equations. Let the *SO(3)-variation facts* be the collection of facts consisting, for every triple of pairwise perpendicular axes through x, of the facts stated by the $SO(3)$-Poisson bracket equations $\{L_i, L_j\} = \epsilon_{ijk} L_k$; where i, j, k take values in the labels of these axes. According to my proposal, what is conveyed about a property by saying that it is represented by an $SO(3)$-Casimir invariant is that the ordinary $SO(3)$-invariance of this property is *fully determined* by the $SO(3)$-variation facts; a determinative relationship which corresponds to the derivability of one collection of Poisson bracket equations from another collection of Poisson bracket equations. This proposal is illustrated in figure (5).

$$\begin{tikzcd} \{f_Q, L_i\} = 0 & Q \text{ is an ordinary } SO(3)\text{-invariant} \\ \{L_i, L_j\} = \epsilon_{ijk} L_k \arrow[u, "\textit{derivation}"'] & SO(3)\text{-variation facts} \arrow[u, "\textit{determination}"'] \end{tikzcd}$$

Figure 5: Illustration of the proposal, according to which the mathematical feature of the $SO(3)$-Casimir invariant function f_Q shown on the left is reflected in a similar feature, shown on the right, of the quantity Q represented by f_Q.

To implement this proposal, we need to do two things: first, we need to explain the nature of $SO(3)$-variation facts; and second, we need to understand the nature of the determinative relationship whose mathematical proxy is the derivability between the relevant mathematical propositions.

3.1 Real Generators

What is the content of Poisson bracket equations such as $\{L_1, L_2\} = L_3$? For every axis i through spatial point x, denote by M_i the i-component of total angular momentum about x. Then our question seems to have the following, straightforward answer. The Poisson bracket $\{L_1, L_2\}$ expresses the definite description 'the quantity whose value at every state equals the rate of change of M_1 under rotations about the 2-axis'; and the equation $\{L_1, L_2\} = L_3$ expresses the non-trivial identity of the satisfier of this definite description and M_3.

But there is a problem with this idea. The Poisson bracket $\{L_1, L_2\}$ prominently mentions L_2, the mathematical representative of M_2; and it does so for an important mathematical reason: L_2 is the generator of the flow in the direction of which we're taking the derivative of L_1. By contrast, the proposed definite description mentions only M_1, but not M_2. This suggests that an important part of the content of $\{L_1, L_2\}$ is missed by the straightforward proposal.

A natural thought is to make the definite description expressed by $\{L_1, L_2\}$ reflect the mathematical content of $\{L_1, L_2\}$ more closely. Mathematically, $\{L_1, L_2\}$ can be paraphrased as 'the function whose value at every phase space point equals the derivative of L_1 along the Hamiltonian flow generated by L_2.' This suggests that the definite description stated by $\{L_1, L_2\}$ is 'the quantity whose value at every state equals the rate of change of M_1 along the family state space curves generated by M_2.' As it stands, however, this is obscure: there is no obvious sense in which a *quantity* counts as a generator of a family of state space curves. The key challenge is thus to give content to this sort of definite description.

The mathematical account reviewed in section (2.2) is suggestive of the following proposal. For a quantity Q to be a generator of the state space curves represented by the Hamiltonian flow of its mathematical representative f_Q is for certain facts about the state space distribution of Q to *determine* facts about the state space curves represented by the Hamiltonian flow of f_Q in a way that mirrors the mathematical determination of this flow by f_Q. I will refer to quantities of this sort as *real generators* of the relevant state space curves.

To implement this proposal, we need to do three things. We need to understand the nature of the facts determined by certain facts about real generators. We need to explain the nature of facts about real generators that determine these facts. And we need to say more about the sense in which certain facts about real generators 'determine' facts of the relevant sort.

What is the nature of the facts determined by certain facts about real generators? I said earlier that Hamiltonian flows on the phase space of an n-particle system capture certain *modal* features of the system. For example, the Hamiltonian flow of L_1 can be used to answer questions such as 'what *would* the system be like if it *were* rigidly rotated by 180° about the 1-axis?' Think of the facts determined by the relevant facts about the real generator M_1 as answers to questions of this kind.

More precisely, let's say that state s' is *rotationally accessible* from state s just in case there is an angle ϕ such that s is related to s' by a rotation about the 1-axis by ϕ. Second, stipulate that a proposition p is *1-possible* at a state s iff p holds at some state rotationally accessible from s, and that p is *1-necessary* at s if p holds at every state rotationally accessible from s. Third, for any two states s and s' related by a rotation about the 1-axis, let the *modal distance* between s and s' be the smallest number $|\phi|$ such that s and s' are related by a rotation by an angle of $\pm\phi$.[31] Finally, let the *1-modal facts* at a state s be the collection of facts about what is 1-possible and 1-necessary at s, together with facts about the modal distances between s and states rotationally accessible from s. The facts determined by the relevant facts about the real generator M_1 are the 1-modal facts at every state.

What is the nature of the facts about a real generator that are supposed to determine these modal facts? Start by observing that the 1-modal facts at every state are characterised by the Hamiltonian flow of L_1: a state s is rotationally accessible from another state s' just in case the phase space points representing s and s' lie on the same Hamiltonian flow line of L_1. This suggests that we can use the mathematical account of the facts that determine a particular Hamiltonian flow line of L_1 as a blueprint for an account of the nature of the facts that determine the 1-modal facts at every state.

According to the mathematical account reviewed in section (2.2), the Hamilto-

[31]To obtain a reparametrisation-invariant measure of the modal distance—a measure invariant under a change from *degrees* to *radians*, for example—an appropriate normalisation factor needs to be added: $\frac{1}{2\pi}|\phi|$ in radians and $\frac{1}{360°}|\phi|$ in degrees. In mathematics jargon, this measure is referred to as the *Haar measure* on the underlying group; cf. (Fuchs and Schweigert, 2003, pp. 375-8).

nian flow of a smooth phase space function is determined by the exterior derivative of this function, i.e. by facts of two sorts: first, facts about which regions of phase space are level surfaces of the function; and second, facts about the differences in value of the function between these level surfaces. And a particular flow line is selected by additionally specifying a particular phase space point—the 'initial value' of this flow line. For example: the exterior derivative of L_1, together with a specific point (q, p), determines which phase space points are related to (q, p) by a rotation about the 1-axis.

The facts about a function f_Q captured by its exterior derivative correspond to facts about the state space distribution of the *physical quantity* Q represented by f_Q, facts I'll refer to as the *level surface facts* about Q: first, facts about the regions of state space in which the quantity has a given constant value; and second, facts about the differences in value of this quantity between those regions. And in close analogy to the mathematical story just described, to infer which states are rotationally accessible from actuality—to infer the actual 1-modal facts—one also needs to know which state one is in; i.e. one needs to know the actual position and momentum facts. This suggests that the range of facts which determine the 1-modal facts are the level surface facts about M_1, together with the position and momentum facts.

Finally: what is the sense in which these two sorts of facts 'determine' the relevant modal facts? A certain metaphysically-minded philosopher might be tempted to deploy the familiar and well-entrenched notions of supervenience or necessitation with regard to necessity of some appropriate type.

These options can be quickly ruled out. Consider the claim that, for the 1-modal facts to be 'determined' by the level surface facts about M_1, together with the position and momentum facts, is for the 1-modal facts to *supervene with metaphysical necessity* on the level surface facts about M_1, together with the position and momentum facts. We have a clear grip on the supervenience of the 1-modal facts on the position and momentum facts: there can be no difference in which states are rotationally accessible without a difference in which possible state is actual. By contrast, if we hold fixed the actual state of the system, i.e. the actual position and momentum facts, then the 1-modal facts are *metaphysically necessary:* given that the actual state of n-particle system is thus-and-such, and given what it is for the system to be 'rigidly rotated' about the 1-axis, facts about which states are rotationally accessible from actuality could not have failed to obtain. This implies that, given the position and momentum facts, the 1-modal facts supervene with metaphysical necessity on any fact whatsoever. But it is not the case that, given the position and momentum facts, the 1-modal facts are 'determined' by any fact whatsoever. So supervenience is inadequate to account for the relevant sense of 'determination'. Similar considerations doom necessitation-based proposals.[32] If there is no type of necessity stronger than metaphysical necessity,[33] then no kind of supervenience or necessitation is adequate to account for what it is for the level surface facts about M_1 to 'determine' the 1-modal facts at every state, given the position and momentum facts.

A natural alternative is *metaphysical grounding*—an explanatory relation that, for present purposes, we may think of as a relation between facts.[34] One key feature of grounding is that 'p metaphysically necessitates q' does not entail 'p grounds q'. The following grounding-based account thus sidesteps the problems of supervenience- and necessitation-based accounts: for M_1 to be the real generator of the state space curves corresponding to rigid rotations about the 1-axis is for the 1-modal facts to be *fully grounded* in the level surface facts about M_1, together with the position and momentum facts.[35]

A grounding-based account captures an important aspect of 'determination'. Whereas the mathematical proxy of the 'determination' relation is *extensional* in the sense that every phase space function co-extensive with f counts as a generator

[32]Given the position and momentum facts, the 1-modal facts are metaphysically necessitated by any fact whatsoever. But it is not the case that, given the position and momentum facts, the 1-modal facts are 'determined' by any fact whatsoever.

[33]Let A-necessity be *stronger than* B-necessity just in case every A-necessary truth is also B-necessary, but not conversely.

[34](Fine, 2001, 2012; Rosen, 2010).

[35]Given the holding of the non-trivial identities stated by the Poisson bracket equations $\{L_i, L_j\} = \epsilon_{ijk} L_k$, there are different, equally correct ways to *state* the relevant grounding claims. For example, the claim that the 1-modal facts are fully grounded in the level surface facts about M_1, together with the position and momentum facts, can be stated just as accurately as the claim that the 1-modal facts are grounded in the level surface facts about *the quantity whose value at every state is equal to the rate of change of M_2 along state space curves generated by M_3*, together with the position and momentum facts. In other words: the relata of grounding involve the relevant generating quantities, but not the definite descriptions by which these quantities can be picked out.

of the Hamiltonian flow of f, we have reason to deny that every quantity state-space co-extensive with the real generator of a family of state space curves is also a real generator of these curves. Even if the total energy of the n-particle system is in fact state-space co-extensive with the quantity equivalent to the sum of kinetic and Newtonian gravitational energy, someone who thinks that the *latter* counts as a real generator of dynamical state space curves would derive wrong claims about the dynamical evolution of the system in the metaphysical possibility in which the system contains electrostatic energy in addition to kinetic and gravitational energy. A grounding-based account reflects this feature of 'determination': if certain modal facts are fully grounded in the level surface facts about a quantity Q, together with the position and momentum facts, then it is generally not the case that, for any quantity Q' state-space coextensive with Q, these modal facts are fully grounded in the level surface facts about Q', together with the position and momentum facts.

Finally, a grounding-based account of 'determination' is a good match for the explanatory connotations of the 'generating' idiom. For example: when say that a device is a *generator* of electricity, we literally mean that it *produces* electricity in the sense that there is an underlying physical process in virtue of which the generator is *causally responsible for* the electricity produced by it. The occurrence of the 'generating' idiom in the present context is therefore suggestive of an account of 'determination' in terms of an appropriate explanatory relation of the same productive species. Since the 'determination' in question is clearly non-causal, grounding again seems like a natural choice: when p grounds q, then q may be thought of as 'produced' by p in a sense that may be paraphrased as 'p is *responsible* for q' or as '*p makes it the case that q*'. Grounding is thus well-placed to account for the productive sense in which the level surface facts about some quantity, together with the position and momentum facts, 'determine' the corresponding modal facts.

The notion of a real generator might seem abstract and unintuitive. But at least one aspect of the proposed explanations of certain modal facts by appeal to the level surface facts about the corresponding real generators should be familiar from everyday modal claims. According to the account just presented, the fact that some proposition p is 1-possible is fully grounded in the level surface facts about M_1, together the position and momentum facts. Now, the level surface facts about

M_1, together with the position and momentum facts, include the fact that all rotationally accessible states agree about M_1, and *a fortiori* they include the fact about the *actual* value of M_1. Therefore, the fact that p is 1-possible is partially grounded in, and thus partially explained by, the fact about the actual value of M_1. In other words, the proposed account entails that an invariant feature of actuality partially explains what is possible, in the relevant sense of 'possible'.

This is reminiscent of more familiar types of modality. When we make claims about what is possible, we always hold certain contextually relevant features of actuality fixed. For example, if you want to know whether it's possible to fit eight people in my car, you are presumably holding fixed the actual size of my car. And the features of actuality we hold fixed in making these claims often seem explanatory of the latter: the fact that it's impossible to fit eight people in my car is explained by the fact that my car has room for no more than five people.

In sum: the notion of a real generator allows us to understand the definite descriptions stated by Poisson brackets, and thus the nature of variation facts. The Poisson bracket $\{L_1, L_2\}$ expresses the definite description 'the quantity whose value at every state is equal to the rate of change of M_1 along state space curves generated by M_2'; and $\{L_1, L_2\} = L_3$ expresses the non-trivial identity of the satisfier of this definite description and M_3.

Moreover, we can now sharpen our understanding of ordinary $SO(3)$-invariance. Let Q be the quantity represented by the function f_Q. The fact that Q is an ordinary $SO(3)$-invariant consists in the conjunction, for every axis i through spatial point x, of the variation facts stated by $\{Q, L_i\} = 0$.

We now have almost all ingredients to state precisely what is conveyed about a property by saying that it is represented by an $SO(3)$-Casimir. The only remaining task is to explain the sense in which the $SO(3)$-variation facts 'determine' the fact that M, the magnitude of angular momentum, is an ordinary $SO(3)$-invariant—a determination which corresponds to the derivability of each of the Poisson bracket equations $\{L^2, L_i\} = 0$ from the Poisson bracket equations $\{L_i, L_j\} = \epsilon_{ijk} L_k$.

3.2 Fine-Grained Invariance

What metaphysical notion is suitable to account for the relevant sense of 'determination'? Our predicament with regard to this question is very similar to that of the previous section, so we can be brief. Supervenience and necessitation are ruled out by the fact that, given what it is for the n-particle system to be 'rigidly rotated' about a point, it is *metaphysically necessary* that magnitudes of vector quantities are $SO(3)$-invariant. Intuitively, that the magnitudes of vectors don't change under rotations follows from *what it is* to rotate a vector. This means that the ordinary $SO(3)$-invariance of any vector quantity supervenes with metaphysical necessity on, and is metaphysically necessitated by, any fact whatsoever.

As before, a natural alternative is metaphysical grounding. For reasons mentioned in the previous section, the following grounding-based account sidesteps the problems of supervenience- and necessitation-based accounts: for the ordinary $SO(3)$-invariance of M to be 'determined' by the $SO(3)$-variation facts is for the fact that M is an ordinary $SO(3)$-invariant to be *fully grounded* in the $SO(3)$-variation facts.[36]

Another advantage of this grounding-based account is that it captures the explanatory characteristics of the relevant sense of 'determination'. Intuitively, facts about how every *component* of a vector changes under a given type of transformation (such as rotations) explain facts about whether the *magnitude* of this vector is invariant under these transformations. In the present context, the fact that the *function* L^2 is an $SO(3)$-constant, i.e. the fact that $\{L^2, L_i\} = 0$ for every axis i through x, is explained by the fact that equations $\{L_i, L_j\} = \epsilon_{ijk} L_k$ hold—i.e. by facts about how the components of $\vec{L}$ change under rotations. This suggests that the ordinary $SO(3)$-invariance of the magnitude of total angular momentum is explained by facts about how total angular momentum components change under rotations—i.e. by the $SO(3)$-variation facts. The grounding-based account of 'determination' vindicates explanatory claims of this sort.

The result of the investigation in this paper is thus the following: what is

[36] Note that the $SO(3)$-variation facts are closed under substitutions of co-denoting definite descriptions. For example, the variation fact stated by the equation obtained from $\{L_1, L_2\} = L_3$ by replacing L_1 with $\{L_2, L_3\}$, L_2 with $\{L_3, L_1\}$, and L_3 with $\{L_1, L_2\}$ is also among the $SO(3)$-variation facts. It follows that the relevant grounding claims are preserved under substitutions of co-denoting definite descriptions of the sort expressed by $SO(3)$-Poisson bracket equations.

conveyed about a quantity Q by saying that its mathematical representative is an $SO(3)$-Casimir invariant is that the ordinary $SO(3)$-invariance of Q is fully grounded in the $SO(3)$-variation facts—a feature I refer to as the *fine-grained $SO(3)$-invariance* of this quantity:

> **Fine-Grained SO(3)-Invariance.** For a quantity Q to be a *fine-grained $SO(3)$-invariant* is for the fact that Q is an ordinary $SO(3)$-invariant to be fully grounded in the $SO(3)$-variation facts.

Fine-grained invariance is the notion we've been looking for. I noted earlier that the ordinary invariance of a physical quantity under a symmetry group is too coarse-grained to account for the link between properties and symmetries: a quantity can be invariant under a group without being 'associated' with that group. By contrast, fine-grained invariance has exactly the right fineness of grain: the fine-grained $SO(3)$-invariants are all and only the quantities 'associated' with $SO(3)$ in the sense identified by physicists.

4 Conclusion

In this paper, I have done two things. First, I reviewed the standard mathematical account of the 'association' between properties and symmetries by introducing the notion of a Casimir invariant. Second, I explained what is conveyed *about a property* by saying that it is represented by a Casimir invariant under some group G. According to the proposal developed in this paper, the answer consists in the notion of a fine-grained G-invariant: that is, the notion of a property whose ordinary G-invariance is fully grounded in the G-variation facts.

The results of this paper carry over *mutatis mutandis* to quantum theory. In that context, the relevant physical quantities are represented by self-adjoint operators on Hilbert space and the relevant families of state space curves by *unitary* flows on that space. Although there is no strict quantum-theoretic analogue of the symplectic form, *Stone's theorem* ensures that every self-adjoint operator generates a unitary flow which preserves this operator.[37] Finally, the role of the Poisson bracket is played by the commutator bracket on the space of self-adjoint operators.

[37] (Stone, 1932).

The account presented in this paper creates conceptual space for a view according to which 'being a fine-grained invariant' under some symmetry group is *essential* to every relevant property of elementary particles (such as mass and spin). In the present context, the idea is that 'being a fine-grained $SO(3)$-invariant' lies in the essence of the magnitude of angular momentum.[38] The result of the present investigation thus promises a novel metaphysics of properties, one which has yet to be articulated. Evaluating whether the link between properties and symmetries calls for a revision of metaphysics along structuralist lines thus has to await further study.

Acknowledgements

I gratefully acknowledge financial support by Princeton University as well as by the *Thomas und Ulla Kolbeck Stiftung*, Germany. I would like to thank David Albert, Gordon Belot, Craig Callender, Adam Caulton, Shai Chester, Neil Dewar, Andreas Ditter, Robert Hirsch, Desmond Hogan, Laurenz Hudetz, Josh Hunt, James Ladyman, Harvey Lederman, Kerry McKenzie, Alex Meehan, Thomas Møller-Nielsen, Jill North, Anna Ross, Pablo Ruiz de Olano, Simon Saunders, Charles Sebens, Ted Sider, David Wallace, James Owen Weatherall, as well as audiences at the University of California, San Diego, at the University of California, Irvine, at the *Direct Empirical Status and the Ontology of Symmetries in Physics* conference at the *Université Catholique de Louvain*, and at the *5th International Summer School in Philosophy of Physics* in Saig, Germany, for helpful comments and discussions. Above all, I owe thanks to Cian Dorr, Hans Halvorson, Boris Kment, Tim Maudlin, and Gideon Rosen.

References

Brading, K. and Castellani, E. (2007). Symmetries and Invariances in Classical Physics. In Butterfield, J. and Earman, J., editors, *Philosophy of Physics Part B*, pages 1331–67. Amsterdam: Elsevier.

Dasgupta, S. (2016). Symmetry as an Epistemic Notion (Twice Over). *British Journal for the Philosophy of Science*, 67(3):837–878.

[38]Modifications to this strategy have to be made for group actions which, unlike $SO(3)$, are associated with more than one fine-grained invariant. I take up this issue in future work.

Fine, K. (2001). The Question of Realism. *Philosophers' Imprint*, 1(1):1–30.

Fine, K. (2012). Guide to Ground. In Correia, F. and Schnieder, B., editors, *Metaphysical Grounding*, pages 37–80. Cambridge University Press.

French, S. (2014). *The Structure of the World: Metaphysics and Representation*. Oxford University Press.

Fuchs, J. and Schweigert, C. (2003). *Symmetries, Lie Algebras and Representations: A Graduate Course for Physicists*. Cambridge Monographs on Mathematical Physics. Cambridge University Press.

Goldstein, H., Poole, C., and Safko, J. (2001). *Classical Mechanics*. Addison Wesley, 3rd edition.

Heisenberg, W. (1976). The nature of elementary particles. *Physics Today*, 29(3):32–39.

Kantorovich, A. (2003). The priority of internal symmetries in particle physics. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 34(4):651–675.

Kochen, S. and Specker, E. (1967). The Problem of Hidden Variables in Quantum Mechanics. *Journal of Mathematics and Mechanics*, 17:59–87.

Ladyman, J. (1998). What is Structural Realism? *Studies in the History and Philosophy of Science Part A*, 29(3):409–424.

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

Lewis, D. K. (1986). *On the Plurality of Worlds*. Blackwell Publishers.

Livanios, V. (2010). Symmetries, Dispositions and Essences. *Philosophical Studies*, 148(2):295–305.

McKenzie, K. (2014). Priority and Particle Physics: Ontic Structural Realism as a Fundamentality Thesis. *British Journal for the Philosophy of Science*, 65(2):353–380.

Ne'eman, Y. and Sternberg, S. (1991). Internal Supersymmetry and Superconnections. In Donato, P., Duval, C., Elhadad, J., and Tuynman, G., editors, *Symplectic Geometry and Mathematical Physics*, pages 326–54. Boston: Birkhäuser.

Noether, E. (1918). Invariante Variationsprobleme. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse*, 1918:235–257.

Penrose, R. (2004). *The Road to Reality.* Jonathan Cape – Random House.

Peskin, M. and Schroeder, D. (1995). *An Introduction to Quantum Field Theory.* Westview Press.

Rosen, G. (2010). Metaphysical Dependence: Grounding and Reduction. In Hale, B. and Hoffmann, A., editors, *Modality: Metaphysics, Logic, and Epistemology*, pages 109–36. Oxford University Press.

Saatsi, J. (2017). Structuralism with and without causation. *Synthese*, 194(7):2255–2271.

Stone, M. H. (1932). On One-Parameter Unitary Groups in Hilbert Space. *Annals of Mathematics*, 33(3):643–648.

Weinberg, S. (1987). Towards the Final Laws of Physics. In Mackenzie, R. and Durst, P., editors, *Elementary Particles and the Laws of Physics – The 1986 Dirac Memorial Lectures*, pages 61–110. Cambridge University Press.

Weinberg, S. (1993). *Dreams of a Final Theory.* New York: Pantheon Books.

Wigner, E. P. (1939). On Unitary Representations of the Inhomogeneous Lorentz Group. *Annals of Mathematics*, 40:149–204.