From the holdings of the Department of Special Collections , Princeton University Library, Princeton, NJ

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyrightlaw. We do not charge any permission or usage fees for the publication of images of material inour collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff.For more information regarding copyright, credit, and citation please review the information onthe Copyright, Credit and Citation Guidelines page of our website.
https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. https://library.princeton.edu/special-collections/ask-us

---

Doctoral Dissertations, 2011-2026 Philosophy

2020

The Theory of Every Thing: Toward a Symmetry-Based Metaphysics of Matter

Schroeren, David

Princeton, NJ : Princeton University

---

http://arks.princeton.edu/ark:/88435/dsp013j333517j

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. https://www.copyright.gov/title17/title17.pdf

The Theory of Every Thing: Toward a Symmetry-Based Metaphysics of Matter

David Philip Berrigan Schroeren[†]

A Dissertation
Presented to the Faculty
of Princeton University
in Candidacy for the Degree
of Doctor of Philosophy

Recommended for Acceptance by
The Department of Philosophy

Advisors:
Hans Halvorson[†]
Gideon Rosen[†]
Boris Kment[†]
Tim Maudlin[*]

June 2020

[†]Princeton University, The Department of Philosophy, 1879 Hall,
Princeton University, Princeton, NJ 08544
[*]New York University, Department of Philosophy, 5 Washington Place,
New York, NY 10003

© by David Philip Berrigan Schroeren, 2020. All rights reserved.

Abstract

In modern physics, almost everything about the material constituents of the world (such as particles) can be inferred from claims about the symmetries of the world. For example, recent breakthroughs in particle physics are almost entirely due to claims about symmetries. This gives rise to an intriguing metaphysical vision according to which the symmetries of the world are fundamental, whereas the material constituents of the world are ontologically derivative of them.

This dissertation develops a general framework for a metaphysics of modern physics that underwrites this vision: an account according to which symmetries correspond to fundamental ontological items and which details how every non-fundamental claim about the material constituents of the world can be accounted for in terms of claims about those items. This account provides a template for an ontology for a wide range of current and future physics and promises to explain the enormous success of symmetry techniques in the recent history of physics.

In the first two chapters, I develop ontological hypotheses that give this metaphysical vision a precise metaphysical underpinning in the relatively simple settings of classical particle mechanics and non-relativistic quantum mechanics, respectively. These views amount to proofs of concept: they show how every relevant claim about material constituents—point particles in the classical setting and wavefunctions in the quantum setting—can be reduced to some claim about symmetries. The third chapter showcases an important instance of the explanatory power and metaphysical utility of symmetry-based metaphysical theorizing. By developing an account of the physical entities that correspond to symmetry groups—the mathematical objects that are used to systematize types of symmetry transformations, such as rotations—this chapter demonstrates how symmetry-fundamentalist ontological theorizing allows us to provide deep, hitherto unavailable explanations of important features of physical properties, such as the fact that quantum-mechanical spin is discrete. The fourth and final chapter develops a perspicuous account of what it is for a physical property to be invariant under a symmetry.

Acknowledgements

I have been incredibly fortunate to work with a committee of stellar philosophers. My advisors complemented each other perfectly and each contributed to the success of my dissertation in innumerable and indispensable ways. Throughout my time at Princeton, Hans Halvorson has gone above and beyond what might reasonably be expected of an advisor: in addition to his invaluable feedback and our inspiring discussions, Hans has been a caring mentor and a steady guide through the daunting waters of a PhD program in philosophy. My progress over the last few years owes much to the way in which Gideon Rosen would invariably turn my oftentimes half-baked ideas into fully-formed philosophical thoughts. It is hard to overstate how much I have learned from Gideon—both about philosophy and about how to be a better philosopher. My dissertation would not be what it is today without Boris Kment. Without fail, his trenchant questioning opened my eyes to critical metaphysical presuppositions and implications of my arguments. Finally, my meetings with Tim Maudlin have been something of a trial by fire for my work. His fierce and relentless criticism always resulted in huge advances in my thinking, and much of what I have achieved in this dissertation I owe to Tim's brilliance, enthusiasm, and generosity.

The graduate community in Princeton's philosophy department has been exceedingly friendly, supportive, and welcoming to me, and I could not have hoped for a better environment in which to take my first steps towards a career in philosophy. Special thanks is due to Daniel Berntson, Ellie Cohen, Ali Fatollahi, Robbie Hirsch, Brendan Kolb, Kyle Landrum, Alex Meehan, Cory Nichols, Sam Preston, Chris Register, and Snow Zhang for helpful discussions, general support, and friendship. It is hard to imagine how I would have made it this far without the help of the philosophy department's administrative workers—especially Anna Faiola, Nancy Groll, Jo Kelly, Jason Robinson, and Joann Zuczek—and I am extremely grateful for their hard work and support throughout my PhD. I gratefully acknowledge financial support from Princeton's Graduate School, the donors to The Bayard Henry, Class of 1876, Graduate Fellowship Fund, as well as Christopher Forbes and Steve Forbes, benefactors of The Malcolm S. Forbes Merit Fellowship. I would also like to thank the Thomas und Ulla Kolbeck Stiftung, and Thomas Kolbeck and Ulla Kolbeck specifically, for their generous support and the faith they have placed in me over the last eight years.

This thesis draws upon and contributes to traditions of philosophy and physics that are, like myself, recent arrivals to Princeton and New Jersey. I would like to acknowledge the intellectual and cultural traditions of the Lenape people who have been dispossessed of the land on which Princeton University stands today—traditions which endure in the Lenape diaspora as well as in the Lenape community still residing on their traditional territory.

Perhaps the most significant intellectual debt I owe outside Princeton is to Cian Dorr. Through one-on-one meetings and various graduate classes at NYU, Cian has had a profound impact on my philosophical thinking and writing, and I hope his influence will be visible throughout this dissertation. I consider it an extremely lucky confluence of circumstances that a significant segment of my recent academic trajectory—spanning the BPhil and the PhD—coincided with Andreas Ditter's. Our friendship has enriched my work and my life in more ways than can properly be expressed here. Many more people deserve my gratitude for feedback, discussions, general support and friendship: David Albert, Claus Beisbart,

Gordon Belot, Alex Blum, Harvey Brown, Shai Chester, Craig Callender, Elena Castellani, Adam Caulton, Eddy Chen, Florian Dahl, Maaneli Derakhshani, Neil Dewar, Sam Fletcher, Steven French, Alexei Grinbaum, Wolfgang Hasse, Meir Hemmo, Mike Hicks, Laurenz Hudetz, Andreas Hüttemann, Josh Hunt, Caspar Jacobs, James Ladyman, Vassilis Livanios, Casey McCoy, Kerry McKenzie, Samuel Meister, Thomas Møller-Nielsen, Fred Muller, Nuria Muñoz, Sebastian Murgueitio, Frédéric Nef, Jill North, Daniel Olson, Sébastien Rivat, Lorenzo Rossi, Pablo Ruiz de Olano, Simon Saunders, Jonathan Schaffer, Chip Sebens, Ted Sider, Trevor Teitel, Barbara Vetter, David Wallace, and Jim Weatherall.

There are many people whose unstinting friendship and support has got me where I am today. Jonathan Berberich, Nils Bergmann, Simon Eschen, and Timo Rößner have been dear friends and loyal companions since high school, and I feel incredibly grateful to have them in my life. I am thankful for the many friendships I made over drinks in the Megaron at Balliol, especially Elwyn Davies, Hannah Uckat, Christopher Jones, Antonia Papenheim, and Carla Roth. I am grateful to Anna Ross, not least for her encouragement and guidance as I began navigating the difficult straits of the academic humanities. Ben Pope has been there for me ever since we started off as neighbors in Jowett; a few years later, we both found ourselves in New York City, extremely online. I am grateful to have gained a treasured friend in Laura Trucco—not even a transcontinental distance has stood in our way. My time in Princeton and New York City would not have been the same without David Amiel, Victor Amoreux, René de Nicolai, Nadja El Kassar, Olivia Gonzalez, Alex Kaiserman, Charles Kollmer, Gabby Paluch, Daniel Putnam, Abdul Rad, Sondre Solstad, Marcus Spiegel, William Stewart, Richard Tuckwell, Max Watkins, and Anna Ziv.

Finally, I would like to thank my parents, Elisabeth and Michael, and my sister Lea. The love and support they have given me is beyond measure; with deep gratitude and appreciation, I dedicate this thesis to them.

Contents

1 Introduction **2**
- 1.1 State-Space Symmetries . . . 8
- 1.2 Varieties of Symmetry . . . 12

2 Symmetry Fundamentalism: A Case Study from Classical Physics **15**
- 2.1 State-Space-First Metaphysics . . . 16
- 2.2 Individualistic Profiles from Primitive States . . . 20

3 Symmetry Fundamentalism in Quantum Mechanics **34**
- 3.1 Symmetry-Fundamentalist Ontology . . . 36
- 3.2 Vector Substantivalism . . . 38
- 3.3 Ray substantivalism . . . 48
- 3.4 A Taxonomy of Ray-Substantivalist Views . . . 50
- 3.5 Discussion . . . 54

4 The Ontology of Symmetry Groups **62**
- 4.1 Groups and Group Actions . . . 65
- 4.2 Group Relationalism . . . 69
  - 4.2.1 Free Actions . . . 72
  - 4.2.2 Non-Free Faithful Actions . . . 73
- 4.3 Discussion . . . 77

5 The Metaphysics of Invariance **81**
- 5.1 Casimir Invariance . . . 86
  - 5.1.1 The Framework . . . 86
  - 5.1.2 Hamiltonian Flows and their Generators . . . 91
  - 5.1.3 SO(3) Poisson brackets . . . 95
  - 5.1.4 Casimir Invariants . . . 98
- 5.2 Fine-Grained Invariance . . . 100
  - 5.2.1 Modality, Variation, and Determination . . . 102
  - 5.2.2 Fine-Grained Invariance . . . 109
- 5.3 Ontic Structural Realism . . . 110

6 Conclusion **113**

Chapter 1

Introduction

In the second half of the 20th century, fundamental physics experienced a major breakthrough: physicists noticed the incredible power of the notion of *symmetry*. The notion of symmetry is familiar: a round dinner plate has a rotational symmetry in the sense that it looks the same if we rotate it. Now, just on the basis of claims about symmetries, physicists were able to infer claims about the nature of physical properties as well as about the nature and behavior of matter—claims that were of unprecedented empirical accuracy. This breakthrough amounts to nothing less than a revolution: *ask not what the world contains but what its symmetries are.*

This paradigm shift towards symmetries gives rise to an intriguing metaphysical vision. When an entity plays a central role in the inferential practice of physics, we are used to entertaining the hypothesis that this entity is *metaphysically* fundamental. The truth of such a hypothesis would be a particularly compelling vindication of this practice: the relevant notion is inferentially significant precisely *because* it has fundamental status and so is involved in the bottom-line story of physical reality that fixes everything else. In this vein, physicists have suggested what I refer to as *symmetry fundamentalism*: the view that symmetries are fundamental aspects of physical reality, whereas the physical entities that we would ordinarily have thought of as the fundamental building blocks of the physical world—such

as elementary particles or fields—are ontologically derivative of these aspects. For example, Steven Weinberg has argued that "symmetries are fundamental"[1] and thus that "at the deepest level, all we find are symmetries and responses to symmetries,"[2] whereas "matter [...] loses its central role in physics."[3] According to Werner Heisenberg, we should "replace the concept of a particle" with "the concept of a fundamental symmetry", and "what we have to look for are not fundamental particles but fundamental symmetries."[4]

Symmetry fundamentalism challenges what is perhaps the default view about the kind of picture of fundamental reality delivered by physics. According to this view, physics describes the world as having fundamental material content: physical objects that *constitute* the world; that count as the *fundamental building blocks* of the world; objects that *compose* the world; entities that at every time instantiate the fundamental properties and relations—properties that figure in the fundamental laws and that enter into causal relations; objects that have a fundamentally spatiotemporal nature, for instance, by virtue of having a spatiotemporal location or a spatiotemporal extension; physical objects whose metaphysical, dynamical and spatiotemporal properties and relations determine and explain the familiar world of our everyday experience. Versions of this assumption are at the heart of a range of influential positions in metaphysics and philosophy of physics, from Democritean atomism via Aristotle's 'material plenum' and David Lewis' 'Humean Mosaic' to the notion of a 'local beable' in the contemporary philosophy of quantum mechanics. But if we follow recent physics, this kind of picture is spectacularly wrong. Fundamental physical reality is not fundamentally material at all: it does not have material constituents at all.

It is far from obvious what picture of fundamental reality is suggested by symmetry fundamentalism. Of course, the negative claim that fundamental reality has no material constituents does not amount to a positive account that details the fundamental structure of

[1](Weinberg, 1987, p. 79).

[2](Weinberg, 1987, p. 80).

[3](Weinberg, 1992, pp. 138-9). Steven Weinberg has confirmed in personal correspondence that these remarks were intended as claims about fundamental physical reality, rather than merely about the epistemic or pragmatic significance of symmetry techniques.

[4](Heisenberg, 1973, p. 273). Cf. Heisenberg (1976, p. 924) and Ne'eman and Sternberg (1991, p. 327)

the physical world. Moreover, the claim that the fundamental physical description features symmetries leaves quite a lot to be desired. Symmetries in physics are introduced as part of the mathematical formulation of physical theories. But symmetry fundamentalism is not the bizarre claim that physical reality literally consists in mathematical objects; it is a thesis about the physical constituents of fundamental reality described in terms of those objects. To give precise content to symmetry fundamentalism, we therefore need to provide a physical ontology: a hypothesis about the ontological inventory of fundamental reality (such as individuals, properties and relations) that makes transparent which physical entities correspond to symmetries and that illuminates how the material constituents of the world—things like particles and fields—are ontologically derivative on these entities.[5]

The goal of this dissertation is to live up to this challenge. I articulate a systematic and unified account that gives symmetry fundamentalism a precise ontological underpinning and promises to explain the theoretical centrality and the empirical success of symmetry techniques in modern physics. I examine symmetry fundamentalism in the fairly simple and mathematically well-defined context of classical particle mechanics and non-relativistic quantum mechanics. The resulting ontological hypotheses function as proofs of concept: they illustrate the nature of symmetry-fundamentalist ontological theorizing in a comparatively straightforward setting and point towards natural extensions to more cutting-edge physical theories. The central component of these hypotheses is what I call *state-space fundamentalism*, the view that the fundamental arena in which the history of the world unfolds is a space which has the structure of a *state space*: phase space in classical particle mechanics and Hilbert space in quantum theory. The thesis that symmetries correspond to fundamental aspects of reality is vindicated by the observation that, in the resulting ontology, mathematical symmetry structure corresponds to a sparse collection of fundamental relations on state space. It is this feature that promises a natural and fairly straightforward extension of the

[5]Note: the claim that physical objects like particles or fields are not fundamental is not tantamount to the claim that fundamental reality does not contain items from the general ontological category of entities that behave syntactically like first-order individuals. As we will see, symmetry fundamentalism is consistent with the thesis that fundamental reality involves entities that are physical objects in this more general sense.

resulting ontological hypotheses to newer physical theories, such as quantum field theory (QFT) and quantum gravity. Both types of theories contain a fundamental notion of a state space, together with certain specific mathematical objects defined on that space: for example, field operators in QFT and area and volume operators in loop quantum gravity.[6] These kinds of mathematical objects have two key features: they correspond to or are definable in terms of symmetries in the relevant sense; and they capture essential physical structure on state space.[7] The main components of the symmetry-fundamentalist accounts I propose therefore have immediate counterparts in the mathematical formulations of some of the most powerful representatives of current and future physics.

Over and above establishing a general framework capable of providing ontologies for cutting-edge physics, symmetry-based metaphysical theorizing also has the power to explain physical facts that have so far been regarded as brute, ungrounded, or metaphysically unexplainable. This dissertation examines one such example: the fact that spin has the following two characteristic features. The first feature is the *discreteness* of spin: there are only countably many values of spin in a given interval—by contrast to quantities such as mass which have as many values in a given interval as there are real numbers in that interval. The second feature is the half-integer-valuedness of spin: values of spin differ by multiples of 1/2. The default position has so far been that this fact has no explanation: spin is typically regarded as a fundamental property whose characteristic features have no explanation in more fundamental terms. The symmetry-fundamentalist account given in this dissertation demonstrates that this default position is incorrect: the discreteness of spin has an illuminating and deep explanation in terms of the physical counterparts of symmetries.

[6] In loop quantum gravity, the fundamental physical entities are mathematically characterized in terms of *spin networks*: graphs whose links are each endowed with a Hilbert space that is acted upon by a unitary representation of the rotation group $SU(2)$—i.e. by a group acting in terms of state-space symmetries. Much like mass and spin in relativistic quantum theory, the fundamental physical quantities of area and volume are defined in terms of the generators of those state-space symmetries; (Rovelli, 2011).

[7] For the same reason, symmetry-fundamentalist ontological theorizing promises a to give a precise metaphysical underpinning to the so-called 'classification of elementary particle kinds' in terms of the irreducible representations of the Poincaré group. This classification is based on a paper by Wigner (1939) and has since become one of the main pillars of the standard model.

Moreover, ontological theorizing that puts symmetries center stage can help elucidate the metaphysical significance of aspects of physical theories that have so far been regarded as mere artefacts of their mathematical formulation. The second example is the notion of invariance. We are familiar with the idea that a property of an object can be *modally* invariant under a symmetry: a round dinner plate *would* look the same if it *were* rotated. But the mathematical formalism of physics appears to draw more fine-grained distinctions between invariant properties. Out of all the properties that are modally invariant under a given class of transformations (such as rotations), a select few are mathematically characterized as bearing a distinctive relationship to those transformations: for example, although the length of linear and angular momentum are both modally invariant under rotations, only angular momentum is mathematically characterized in this way.[8] This more fine-grained mathematical characterization of certain modal invariants of invariance is at the heart of a range of claims concerning the relative ontological priority of symmetries and the matter content of the world, such as particles—claims that are expressed by physicists and philosophers in equal measure. By putting mathematical symmetries front and center as inspirations for ontological theorizing, this dissertation demonstrates how this fine-grained mathematical characterization of physical properties can be used to precisely articulate a novel account of property invariance that underwrites the relevant ontological and physical distinctions between modal invariants.

But the interest of the ontological hypotheses for classical and quantum mechanics I articulate go beyond both their applications to more cutting-edge physical theories and their explanatory merits; they are novel and interesting metaphysical views in their own right. Most importantly, a symmetry-fundamentalist metaphysics represents a new option for the ontology of non-relativistic quantum mechanics. Moreover, a symmetry-fundamentalist ontology for classical particle mechanics describes physical reality without individualistic (particulate) content and therefore represents a systematic and precise alternative to the standard individualistic ontology for the theory.

[8]The relevant mathematical notion is that of a *Casimir invariant*.

The questions I am investigating in this dissertation have also been a key inspiration for *ontic structural realism*, an influential intellectual tradition in the philosophy of science according to which physical reality is fundamentally 'structural' in a way that is to be explicated at least in part by appeal to symmetries.[9] A discussion that does justice to the rich and vibrant structuralist literature on the topic is beyond the scope of this work, although chapter 5 will briefly expand on the relationship between ontic structuralism and the account I develop in this dissetation. For present purposes, I want to point out two things. First of all, the approach I will pursue does not respect the constraint, expressed by some ontic structuralists as well as other metaphysicians, that an account of fundamental physical reality must not involve first-order individuals.[10] Second, the ontic structuralist project has been quite controversial: for example, a common complaint is that the operative notion of 'structure' has not been satisfactorily clarified.[11] Although I am in agreement with ontic structural realists about the metaphysical importance of symmetries in physics, in this paper I will explore whether the relevant symmetry-based claims can be put on a sound ontological footing without committing to the controversial structural realist assumptions. This is a project that should be of interest even to those who favor an ontic structuralist approach.

This dissertation is structured as follows. In the remainder of this chapter, I introduce some of the basic notions that will be used throughout this work. In particular, I will introduce the notion of symmetry that is operative in the claim that symmetries are fundamental aspects of reality: the notion of a *state-space symmetry*. I will also introduce the notion of a state space and explain the relationship between state-space symmetries and other notions of symmetry, such as dynamical and space-time symmetries. Chapters 2 and 3 develop a symmetry-fundamentalist ontology for classical particle mechanics and non-relativistic quantum mechanics, respectively. Chapter 4 details an account of the ontological counterparts

[9](Ladyman, 1998; Castellani, 1998b; Brading and Castellani, 2003; Kantorovich, 2003; Ladyman et al., 2007; Roberts, 2011; French, 2014; McKenzie, 2014a; Brading et al., 2017).

[10](French and Ladyman, 2003; Ladyman et al., 2007; French, 2010). Other metaphysicians who have developed proposals for dispensing with first-order individuals include Shamik Dasgupta (2009) and Andrew Bacon (2019).

[11]For example, (Dorr, 2010; Wolff, 2011; Arenhart and Bueno, 2015); cf. (McKenzie, 2017).

of symmetry groups; an account that provides the basis for explaining essential characteristics of physical properties such as spin in terms of underlying symmetries of the world. Chapter 5 articulates a fine-grained notion of property invariance under a symmetry that applies to some but not all of the modal invariants under that symmetry. Chapter 6 contains some concluding remarks.

1.1 State-Space Symmetries

Symmetry is a notion of which we have an intuitive grasp: there is a natural sense in which everyday objects (like a round dinner plate) can have rotational symmetries, in the sense that the object looks the same when rotated. The basic idea is that a symmetry of an object corresponds to ways of changing the object that preserve certain salient features of that object. This idea is also reflected in the physical notion of symmetry, albeit in a more abstract form: in this setting, the kind of objects whose distinctive properties are invariant under symmetries are *mathematical spaces*, and symmetries are mathematical functions (also known as *transformations*) on those spaces that preserve those properties.

Here is an example. We are familiar with the idea that Euclidean space has symmetries, such as rigid spatial shifts and rigid rotations. These transformations count as symmetries because they preserve what one might think of as the characteristic structural feature of Euclidean space: the *Euclidean metric*. Symmetries of Euclidean space are transformations that preserve this structure, also known as *Euclidean isometries*.

This example illustrates something important: abstractly, a symmetry of a structured mathematical space is always defined as a transformation of that space which preserves the characteristic structure of that space.[12] The symmetries that are involved in symmetry fundamentalism are symmetries in this more abstract sense: they are symmetries of whatever mathematical space corresponds to the *state space* of the world according to the theory in

[12]In this way, we can think of structured mathematical spaces as individuated by their symmetries along the lines of the *Erlanger Programm* (Klein, 1893, 1910); cf. (Wallace, 2019) for a more recent statement of this idea.

question. Symmetries of this sort are known as *state-space symmetries.*[13]

I will explain shortly why state-space symmetry is the notion of symmetry operative in symmetry fundamentalism. For now, let me tell you what a state-space symmetry is. First of all, a proposition is the *instantaneous state* of a world at a time iff it is a maximally specific and exhaustive description of the values of the variable fundamental physical quantities at that time: that is, the physical quantities whose values possibly vary between instants of time and whose values at some instant of time, taken together, determine the values of all other physical quantities at that instant. The fact that states are exhaustive and maximally specific entails that states are incompatible: the world cannot be in more than one state at any given instant of time. However, the world is necessarily in *some* state at every instant of time. States therefore have some similarities to possible worlds.[14]

For example, consider classical (Newtonian) N-particle mechanics: the theory that describes the temporal evolution of N point particles in three-dimensional Euclidean space according to Newton's laws of motion. The default ontology for this theory is what I call *individualism*: the view that fundamental reality consists in point particles instantiating the values of certain fundamental quantities. A common choice for these quantities is position and momentum. Figure (1.1) is a schematic illustration of a world with two particles.

State space is just the space of all instantaneous states of the world: in the case of classical particle mechanics, the space of all possible combinations of positions and momenta of every particle. In this setting, the mathematical spaces used to characterize state spaces are called *phase spaces*, illustrated in figure (1.2). In non-relativistic quantum mechanics, the state space of a physical system is characterized by *Hilbert space*: a complex inner product space whose elements are also referred to as 'kets'.

This brings us to the notion of a state space symmetry. Just as a transformation of

[13]The notion of a state-space symmetry is well-entrenched in physics and are routinely mentioned in textbooks; for example, (Strocchi, 2005; Moretti, 2013). In physics, state-space symmetries are sometimes referred to as *Wigner symmetries*, reflecting the seminal work on the topic by Eugene Wigner (1931, 1968). State-space symmetries have also been recognized as a distinctive kind of symmetry in the philosophical literature—for example, by John Earman (2003, p. 341) and Simon Kochen (1996, p. 242).

[14]For more on the connection between states and possible worlds, see (Williamson, 2016).

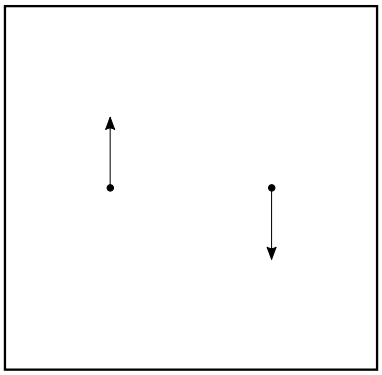


Figure 1.1: Schematic illustration of a two-particle world at some instant of time.

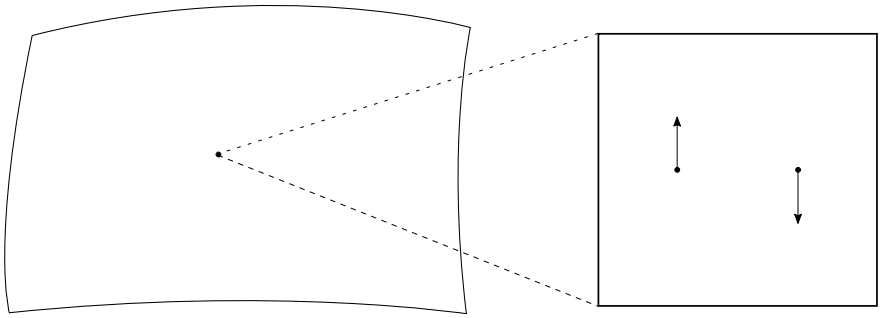


Figure 1.2: Each point in phase space (shown on the left) corresponds to a situation as shown on the right. The state space of a two-particle world contains all possible situations of this sort.

Euclidean space is a symmetry of Euclidean space iff it preserves the characteristic structural feature of Euclidean space—the Euclidean metric—a transformation of the mathematical surrogate of state space is a *state-space symmetry* iff it preserves the characteristic structural feature of this mathematical space: the feature that stands to state space as the Euclidean metric stands to Euclidean space. In the classical setting, the characteristic structural feature is the symplectic form. To a first approximation, you can think of the symplectic form as a mathematical object that determines the volumes of phase-space regions, and of phase-space symmetries as transformations that preserve the volumes of such regions, as illustrated in figure (1.3).[15] In the quantum setting, the characteristic structural feature is the inner product—a function you can roughly think of as an assignment of lengths to vectors—and

[15]The reason why this is only a first approximation is that, although every transformation that preserves the symplectic form also preserves the volumes of phase space regions, the converse is not true.

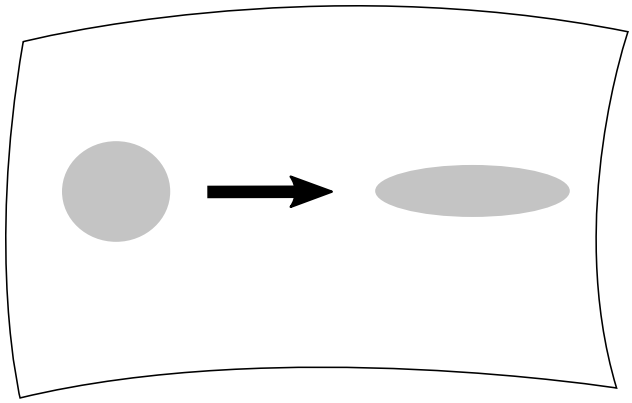


Figure 1.3: An illustration of phase space. The map that corresponds to the black arrow is a symmetry of phase space only if it preserves the volumes of regions in phase space. This means, in particular, that this map is a symmetry only if the two shaded regions are of equal volume.

state-space symmetries are also known as *unitary transformations.*[16] (I will provide more detail about the quantum setting in chapter 3.) An example of state-space symmetry in classical particle mechanics are shifts of individual particles in position and momentum, as illustrated in figure (1.4).

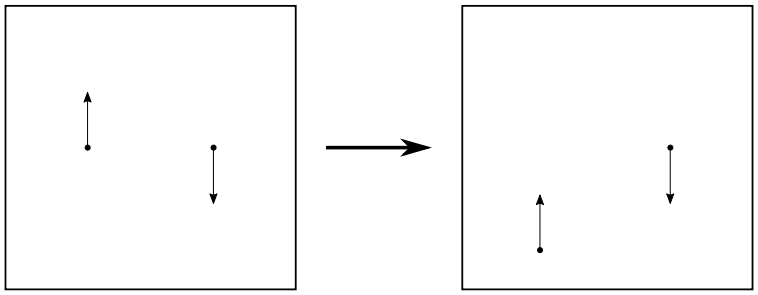


Figure 1.4: The state on the left and the state on the right disagree only about the position of one particle. The arrow indicates the transformation that takes the state on the left to the state on the right. If this transformation preserves the symplectic form (and so the volumes of phase space regions), it is a symmetry of phase space.

[16]Strictly speaking, *anti-unitary* transformations count as symmetries, too; though they play a much less significant role in physics than unitary transformations; cf. (Wigner, 1931; Weinberg, 1995).

1.2 Varieties of Symmetry

A state-space symmetry is not the same thing as a dynamical symmetry: the latter are state-space symmetries that also preserve the laws. More precisely: a state-space symmetry T is a dynamical symmetry iff for every dynamically allowed trajectory h, there is another dynamically allowed trajectory h' such that every point on h gets mapped by T to a unique point on h'. So, every dynamical symmetry is a state-space symmetry, but not conversely.

I should therefore explain why it is the notion of a state-space symmetry that is operative in symmetry fundamentalism, rather than the notion of a dynamical symmetry. The main reason is that, whether a given state-space symmetry is also a dynamical symmetry plays no role in the relevant symmetry-based inferences; and if a dynamical symmetry figures in such an inference, it does so only in virtue of the fact that it is a state-space symmetry. Perhaps the best illustrations of this can be found in the following two examples from quantum theory. First, the only symmetry-related premise in the characterization of spin in non-relativistic quantum mechanics is that rotations are state-space symmetries.[17] Similarly, the only symmetry-related premise in the classification of elementary particle kinds in terms of mass and spin is that relativistic boosts, rotations, and shifts (also known as *Poincaré transformations*) are state-space symmetries.[18] Both are paradigmatic examples of symmetry-based inferences that motivate symmetry fundamentalism; but neither requires an assumption about whether the state-space symmetries in question also count as dynamical symmetries.

Moreover, note that the notion of a state-space symmetry is also distinct from the notion of a *spacetime symmetry*: roughly, a symmetry of the physical space that contains the physical constituents of the world, such as particles, fields, tables and chairs. We already encountered one kind of spatial symmetry: the symmetries of Euclidean space. Other spaces that are physical spaces or spacetimes in this sense are Minkowski space or (more generally)

[17] For example, the fact that spin is discrete follows from the Peter-Weyl theorem on the assumption that rotations are state-space symmetries (i.e. unitary transformations); this is the only symmetry-related assumption required by the theorem (Folland, 2015, p. 136).

[18] This classification is due to Eugene Wigner (1939).

Pseudo-Riemannian manifolds.

To be sure: there are special physical systems whose state-space, dynamical, and spacetime symmetries coincide. For example, consider a classical Newtonian particle world that contains exactly one particle in three-dimensional Euclidean space and suppose that its dynamics are governed by a Hamiltonian of the form $H = \frac{p^2}{2m} + V(x)$, where $V(x)$ is a spherically symmetric potential. In this world, rotations are symmetries of all three kinds: rotations are spatial symmetries (since they preserve the Euclidean metric); they are state-space symmetries (since they preserve symplectic form); and they are dynamical symmetries (because they leave H invariant).

These kinds of assumptions are often tacit in standard presentations of mathematical connections on which the symmetry-based metaphysical theses are based.[19] However, these assumptions do not hold in most cases of physical interest. For a classical n-particle system, one fairly generic type of state-space symmetries are rotations of an individual particle that holds fixed all other particles. But rotations of individual particles do not generally count as a dynamical symmetries: the most generic n-particle Hamiltonians are not invariant under such transformations. (Intuitively, this is witnessed by the fact that there are many ways of rotating an individual particle that result in empirically distinguishable situations.)

Similarly, spatial rotations of individual particles in an n-particle system cannot be subsumed under spatial symmetries of Euclidean space. A rotation of Euclidean space about a fixed point can be thought of as a map that assigns to each point of Euclidean space its image under that rotation; rotations are *permutations* of points in Euclidean space. Think of a rotation of Euclidean space containing n particles as a permutation of which locations are occupied by which of the n particles. Since rotations preserve the Euclidean metric, these permutations do not change particle distances. By contrast, a rotation of an individual particle changes the spatial location of that specific particle while holding the spatial locations of other particles fixed; so, rotations shift that particle both relative to the other particles

[19] For example, (Ballentine, 2000, pp. 77).

and relative to the points of the underlying Euclidean space. *A fortiori*, rotations that are state-space symmetries generally are not equivalent to permutations of points of Euclidean space. Rotations as state-space symmetries are equivalent to rotations of Euclidean space only in the special case of rigid, uniform rotations of all n particles, or in the special case of single-particle worlds. To illustrate: in a world containing only a single classical particle in Euclidean three-space, changing the spatial location of that particle by a rotation about some axis (a state-space symmetry) is tantamount to a permutation of the points of Euclidean space (a spatial symmetry). This is not true in general.

The upshot is twofold: first, although state-space symmetries have systematic relationships with dynamical and space-time symmetries in certain special cases, the notions are distinct. Second, even in the special cases where these three notions coincide, the fact that a given transformation figures in the relevant symmetry-based claims owes only to its status as a state-space symmetry.

Chapter 2

Symmetry Fundamentalism: A Case Study from Classical Physics

The goal of this section is to develop a specific instance of symmetry fundamentalism in the context of classical particle mechanics: the theory whose standard, *individualistic* ontology consists of point particles evolving according to Newton's equations of motion. Taking inspiration from the symmetry structure of the theory, I articulate an alternative ontology that vindicates symmetry fundamentalism. But this chapter is not primarily intended as a contribution to the metaphysics of classical mechanics. Rather, the case study illustrates symmetry-fundamentalist ontological theorizing by appeal to a familiar example and serves as a guide for the development of corresponding ontological proposals for quantum theory. This is a task I take up in chapter 3, in which I articulate a symmetry-fundamentalist ontology for non-relativistic quantum mechanics.

My goal in this chapter is not to give a systematic defense of symmetry fundamentalism, but rather to develop a precise account of what it would be for symmetry fundamentalism to be true at a world consisting of classical point particles. In other words: the aim is to detail a classical-mechanical world whose constituent particles are ontologically derivative on the fundamental ontological counterparts of symmetries. This is a substantive challenge: when

we are told that entities of one kind are ontologically derivative of entities of another kind, we expect an account that details how every relevant proposition about the former is entailed by some proposition about the latter. One of the main tasks of this chapter is therefore to show that, for every proposition about particles statable in the individualistic theory, there is some fundamental proposition about the physical counterparts of symmetries that entails it. Indeed, I will show something even stronger: that every proposition about particles statable in the individualistic theory is *equivalent* to some fundamental proposition.

In the classical setting, the old, materialistic theory is just the individualistic ontology: classical Newtonian particle mechanics that is fundamentally about n point particles in Euclidean space. In the quantum setting, the relevant family of materialistic theories consists in those theories whose fundamental ontology involves the physical field characterized by the wavefunction. More on this in chapter 3.

2.1 State-Space-First Metaphysics

Call a proposition the *complete individualistic profile* of a classical particle world at some instant of time iff it fixes the positions and momenta of all particles at that time. The complete individualistic profile at some time entails the values of all other physical quantities at that time. According to the individualistic metaphysics, instantaneous states are complete individualistic profiles. In other words: for individualists, state space is nothing over and above possible complete individualistic profiles: for the world to be in some state at some instant of time *just is* for some complete individualistic profile to be true.

But there is an alternative view of state space, one that is key to the symmetry-fundamentalist ontology I'm about to develop. According to this *state-space-first* view, there are fundamental entities (called *primitive states*) and some appropriate fundamental notion of 'actuality' which satisfies the constraint that, necessarily, for every instant of time there is exactly one primitive state that is actual at that instant. Primitive states have

the structure of the individualistic state space and are in one-to-one correspondence with complete individualistic profiles. Although the state-space-first metaphysician agrees with the individualist about which complete individualistic profile is true at any given instant of time, they deny that primitive states are complete individualistic profiles: on the state-space-first view, particles are ontologically derivative and so do not figure in the bottom-line story of physical reality. Primitive states are therefore quite different from their individualistic counterparts: as opposed to individualistic states, when two primitive states are distinct, this is not so by virtue of disagreeing about any particular particle; they are primitively distinct.[1] This means that, on the state-space-first view, the constraint that no more than one state can be actual at any given instant of time is quite different in nature than it is on the individualistic view. On the individualistic picture, the fact that at most one complete individualistic profile can be true at any given instant of time is a logical consequence of the fact that individualistic profiles are maximally specific and exhaustive, and therefore mutually exclusive. Since the state-space-first metaphysician regards states as primitive entities that are distinct from complete individualistic profiles, they must implement this constraint as a basic law governing the actuality relation.

Since the state-space-first view entails that particles are ontologically derivative, the main challenge for this view is to show that every complete individualistic profile is entailed by some fundamental proposition about primitive states.

It is clear that the requisite fundamental propositions cannot *just* say which primitive states are actual. Recall: according to the individualist, the proposition that some individualistic state is actual at some instant of time determines the complete individualistic profile at that instant. Since primitive states do not have individualistic structure, the state-space-first

[1](Bacon, 2019) develops a version of state-space-first metaphysics, albeit for somewhat different philosophical purposes: he deploys state-space theoretic reasoning to illustrate a general strategy aimed at the more ambitious goal of articulating a purely qualitative account of the fundamental—one that dispenses with individuals in a broader logical sense of first-order individuals. By contrast, the individualism to which state-space-first metaphysics is opposed is the view that fundamental physical reality consists in individuals of a more specific, physical sort: objects like elementary particles. My kind of state-space-first metaphysics is consistent with the more general kind of individualism that is the target of Bacon's work. The view proposed in (North, 2009) merely entails that state space is a real physical space, but not that it is fundamental.

metaphysician must deny the corresponding claim: the proposition that some primitive state is actual at some instant of time does not determine which complete individualistic profile is true at that time. The state-space-first picture requires additional fundamental entities whose pattern is detailed by the relevant propositions, such as properties of and relations among states. So, whereas individualistic states determine complete individualistic profiles *in isolation*, primitive states determine these profiles only *collectively*. While individualistic states may have some resemblance to possible worlds, primitive states clearly do not. I say more about the modal character of primitive states at the end of this chapter.

To identify the physical entities whose pattern is detailed in the relevant fundamental propositions, we need to specify the state-space-first view in more detail. An immediate question concerns the ontological category of primitive states. This is a subtle choice point for the state-space-first approach. One option is *state-space substantivalism*, the view that primitive states are first-order individuals. In addition to primitive states, this view posits *instants of time*: first-order individuals that bear a fundamental dyadic *actuality* relation to states. This relation can be paraphrased as "... is actual at..." and satisfies the constraint that, necessarily, for every instant of time t there is some state s such that s is actual at t.[2] An alternative view conceives of primitive states as *dyadic relations* between instants of time and a single first-order individual—'the world'—such that necessarily, for every instant of time, there is some relation that the world bears to that instant.[3] Another alternative regards primitive states as *monadic properties* of instants of time such that necessarily, for every instant of time there is some such property instantiated by it.[4] Although I focus on the substantivalist option in this chapter, my main conclusions do not depend on this choice.

[2]The explicit quantification over instants of time means that state-space substantivalism has a distinctively B-theoretic flavor. It is fairly straightforward to articulate an A-theoretic version of the view—a version which does not quantify over instants of time—by replacing the actuality relation with a monadic property '...is actual', and introducing primitive propositional tense operators such as 'it will always be the case that...', 'it will at some time be the case that...' and so on.

[3]The A-theoretic version of this view regards primitive states as *monadic properties* of the world such that necessarily, the is always some such property instantiated by the world.

[4]The A-theoretic version of this view conceives of states as zero-place properties: i.e. *propositions* such that necessarily, there is always some such proposition that is true.

The core of the view I propose is that primitive states are fundamental physical entities that figure in the fundamental propositions that entail complete individualistic profiles. This view is indifferent to whether states are conceived of as first-order individuals, first-order relations between the world and instants of time, or as monadic properties of instants of time.[5]

In addition to states, times, and the actuality relation, we also need to posit properties and relations that confer symplectic structure on primitive states and temporal structure on instants of time. I will not say more about these ontological primitives here, except for the following. First, the literature contains successful proposals for an intrinsic account of temporal structure.[6] Second, an intrinsic account of the properties and relations that confer symplectic structure is tantamount to the nominalization of symplectic geometry. This project is beyond the scope of this dissertation; for present purposes, it will suffice to refer to the symplectic structure in mathematical terms.

A few paragraphs ago, I noted that the complete individualistic profile at an instant of time is determined by a proposition that does not just say which primitive state is actual at that time, but that also details the pattern in certain properties of and relations among primitive states. Are the properties and relations that confer symplectic structure sufficient for this task?

The answer is no. State spaces of classical particle worlds can agree about the entities that confer symplectic structure while disagreeing about particle number. A twelve-dimensional phase space can be used to characterize the state space of *two* particles in *three* spatial dimensions as much as it can be used to describe the state space of *three* particles in *two* spatial dimensions. Complete individualistic profiles aren't determined by symplectic structure alone; additional physical structure is needed. As I will explain shortly, central among these primitives are relations among states that correspond to state-space symmetries.

In this section, I contrasted individualism—the default ontology for classical particle

[5]See (Bacon, 2019) for a discussion of some foundational issues raised by this choice point.

[6]For example, Arntzenius and Dorr (2012, pp. 223) proceed in terms of temporal betweenness and temporal congruence; cf. Field (1980, Chpt. 6).

mechanics—with the state-space-first approach. The next task is to introduce the ontological primitives whose patterns determine complete individualistic profiles. There will be no more than four primitives, two of which correspond to state-space symmetries. This will demonstrate that state-space-first metaphysics is an instance of symmetry fundamentalism.

2.2 Individualistic Profiles from Primitive States

I will introduce the relevant primitives by their individualistic paraphrases; but it is important to emphasize that these individualistic characterisations are not metaphysical definitions: as primitives, these items cannot be characterized or defined in more basic terms. Any individualistic characterisation is a mere conceptual crutch that we must ultimately kick away to be replaced by a full account of individualistic talk in state-space-first terms. This is a familiar aspect of metaphysical theorizing: when making unfamiliar metaphysical posits, we need to defend them by presenting a theory in which they figure and by showing how that theory is inferentially and explanatorily connected to notions that we already understand. This is precisely what I will do momentarily: I will explain how the ontological primitives I introduce figure in the fundamental propositions that entail complete individualistic profiles.

I already mentioned some of the fundamental posits of state-space-first metaphysics. On the state-space substantivalist version of this approach, we posit states, instants of time, the actuality relation, as well as symplectic and temporal structure. (For ease of exposition, in this section I will refer to primitive states simply as 'states'.) There are four additional primitives; here are the first three:[7]

> **1. Position-Almost-Sameness.** A dyadic relation. Relates two states iff they agree about the momenta of all particles and about the positions of all but one particle.
>
> **2. Momentum-Almost-Sameness.** A dyadic relation. Relates two states iff

[7]The following is inspired by (Dorr, 2018).

they agree about the positions of all particles and about the momenta of all but one particle.

3. Degeneracy. A monadic property. Picks out a *6*-dimensional subspace of state space consisting of states at which all particles have the same position and the same momentum.

Before we can introduce the fourth primitive, we need to do some work using the other primitives. There are four steps.

First, position-almost-sameness picks out a range of subspaces such that any two elements are position-almost-same; similarly for momentum-almost-sameness. Call these subspaces the *position-almost-same subspaces* and the *momentum-almost-same subspaces*, respectively.

Second, observe that the points in a given position- or momentum- almost-same subspace differ about the position or momentum *of the same particle*, respectively. For example, the reason is that any two states y, z can be position almost-same to some state x without being position-almost same to each other. The fact that x, y and x, z are position-almost-same, respectively, implies the position-almost-sameness of y and z only if the differences between x, y and x, z concern the same particle.

The position- and momentum-almost-same subspaces are characterized in terms of certain specific state-space symmetries: the position and momentum shift flows, i.e. the transformations that on the individualistic picture correspond to shifts of individual particles in position or momentum. Recall: a phase-space flow is a family of curves such that every phase space point is intersected exactly once by exactly one curve in this family. For any triple of spatial coordinate axes, there is a triple of phase space flows that corresponds to shifting some specific particle along these three axes. Of course, there is a continuum infinity of such triples—one for every coordinate system—each of which is equally well-suited for describing the almost-same subspaces. This is as it should be: differences between coordinate systems are a paradigm instance of mathematical differences without corresponding metaphysical distinctions.

Here is the third step. We can define physical counterparts of level surfaces of coordinate functions in terms of almost-same subspaces. The *level surfaces* of a coordinate function on phase space are regions of phase space in which that function is of constant value. Level surfaces may be familiar from hiking maps where they are known as *contour lines*: regions of constant height. See figure (2.1) for an illustration.

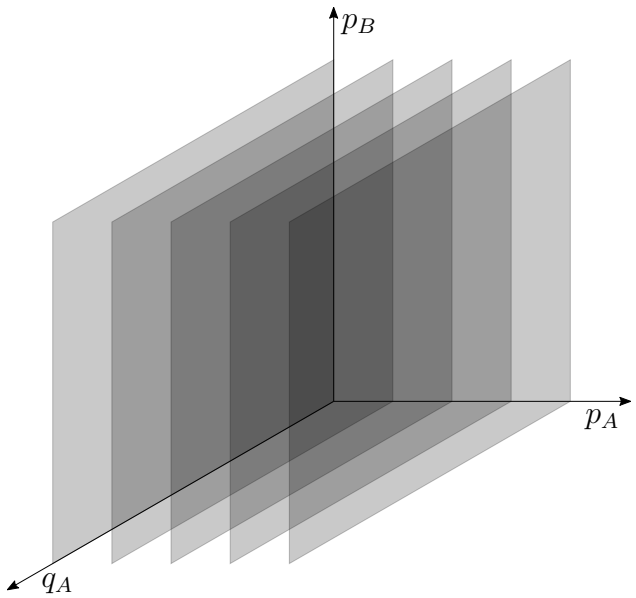


Figure 2.1: An illustration of the phase space of two particles (here labeled by A and B) in one spatial dimension. Axes are labeled by the coordinate functions for the position of particle A (q_A), the momentum of A (p_A), and the momentum of B (p_B); the axis corresponding to the position of B (q_B) is suppressed. The shaded regions are segments of the level surfaces of p_A. (More precisely: the shaded regions are the intersections of the three-dimensional level surfaces of p_A with some specific three-dimensional level surface of q_B).

The physical counterparts of level surfaces are what I refer to as *position* and *momentum foliations*. We can give them the following individualistic characterization: any two states on a leaf of a position foliation agree about the position of some particle while possibly disagreeing about the positions of all other particles and about the momenta of all particles; similarly, any two states on a leaf of a momentum foliation agree about the momentum of some particle while possibly disagreeing about the momenta of all other particles and about the positions of all particles. Position and momentum foliations are equivalence relations with leaves as equivalence classes; and two states x, y are on the same leaf of (say) a position

foliation Q iff Qxy.

Position and momentum foliations are mathematically represented by the level surfaces of phase-space coordinate functions. The subspace of phase space that corresponds to a given leaf of a position foliation is the subspace such that, for any two phase space points s, s' in this subspace, for some particle, for every triple of spatial coordinate functions of that particle, and for some triple of level surfaces of these functions, s lies on those level surfaces iff s' does. Differently put: any two points in a subspace of phase space that corresponds to a given leaf of a position foliation agree about the level surfaces of all spatial coordinate functions of some particle.

The goal is now to replace the individualistic characterization of position and momentum foliations with a definition in terms of position- and momentum-almost-sameness. This definition involves the symplectic structure of state space. As I said before, for present purposes it will be sufficient to refer to this structure in mathematical terms. The key thing to note, then, is that position and momentum foliations can be defined in terms of the position- and momentum-almost-sameness relations at the level of mathematical description: that is, the level surfaces of a given coordinate function can be defined in terms of a specific phase space flow.

The mathematical details of this definition are quite involved and need not concern us here.[8] It will suffice to note that the definition is based on four mathematical observations:

[8]The details go as follows. Let (M, ω) be a symplectic manifold, p a smooth real-valued function on M, $exp_x : T_xM \to M$, $v \mapsto exp_x(v)$ the exponential map and d the exterior derivative. The level surfaces of p are encoded in the exterior derivative dp of p; and dp uniquely corresponds to the phase space flow

$$\begin{aligned} &\gamma : M \times \mathbb{R} \to M, \\ &\quad (x, t) \mapsto \gamma(x, t) = exp_x(t\omega^{-1}(dp)|_x) \end{aligned} \tag{2.1}$$

where $\omega^{-1}(dp)$ is the Hamiltonian vector field of p. Since $\omega^{-1}(dp, dp) = 0$, i.e. the rate of change of p under γ vanishes, γ is confined to the level surfaces of p. Moreover, there is a smooth function q (unique up to an arbitrary additive constant) such that $\omega^{-1}(dq, dp) = 1$. If p is the momentum coordinate function of some particle along some axis, then q is the corresponding position coordinate. For the purposes of ontological theorizing, we can read the identity in (1) as stating a definition of dp (corresponding to an aspect of a momentum foliation) in terms of γ (corresponding to an aspect of a position-almost-same subspace).

To be sure, the identity in (1) by itself does not require this reading: there is nothing in the mathematics that suggests an asymmetry between the two sides of the equality sign. However, we are presently concerned with *ontological theorizing* about physics. As I understand this project here, it consists in the articulation

first, every vector field has a flow; for example, figure (2.2) shows the gradient vector field of a momentum coordinate function. Its flow takes us along paths of greatest increase of this function. Second, there is a one-to-one correspondence between the level surfaces of this

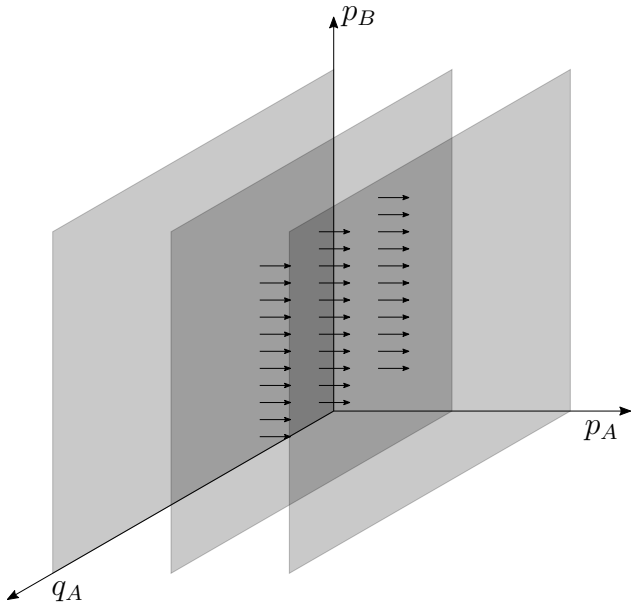


Figure 2.2: The gradient of p_A is everywhere at a right angle to the level surfaces of p_A and its length at a point is proportional to the rate of change of p_A around that point.

coordinate function and its gradient: the level surfaces uniquely determine the gradient and vice versa. Third, for technical reasons, phase spaces generally do not admit of a gradient vector field, but instead of what's referred to as a *Hamiltonian vector field*: a vector field that differs from the gradient in that it does not point in the direction of greatest increase of a function but rather *along the level surfaces* of that function, as illustrated in figure (2.3).[9] But Hamiltonian vector fields are similar to gradients in the following respect: for every Hamiltonian vector field there is some unique collection of level surfaces and a unique flow to which it corresponds. Fourth, the level surfaces of the *position* coordinate of some particle in some direction uniquely correspond to the Hamiltonian vector field whose flows are shifts

of hypotheses about physical reality in a way that may involve drawing metaphysical distinctions between mathematically equivalent entities; cf. (Maudlin, 2018).

[9]The reason is that phase spaces generally do not admit a Riemannian metric, a mathematical entity that links the level surfaces of a function to gradient vector fields. Instead, as already mentioned, phase spaces admit of a symplectic form—an entity which links the level surfaces of any smooth real-valued function with some specific Hamiltonian vector field.

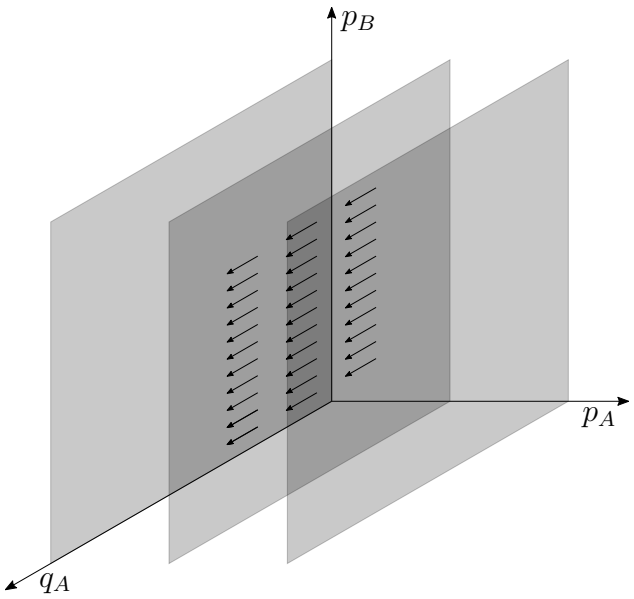


Figure 2.3: The Hamiltonian vector field of p_A everywhere points along the level surfaces of p_A. Its length at a point is proportional to the rate of change of p_A around that point.

in the *momentum* coordinate of that particle in that direction; similarly, the level surfaces of the *momentum* coordinate of some particle in some direction uniquely correspond to the Hamiltonian vector field whose flows are shifts in the *position* coordinate of that particle in that direction. This is illustrated in figure (2.4).

The result is the following: at the level of mathematical description, position foliations are definable in terms of momentum-almost-sameness, and momentum foliations are definable in terms of position-almost-sameness—in both cases, together with the symplectic structure of state space.[10]

The fourth step is to make use of the degeneracy property. The idea is that degenerate states serve to define surrogates for position and momentum values. We begin with defining the relations of *hereditary position-/momentum almost-sameness* as the transitive closures of the position-/momentum-almost-sameness relations. In other words: any two hereditarily position-almost-same states agree about the momenta of all particles but possibly disagree

[10]The definitions of position-foliations in terms of momentum-almost-sameness and of momentum-foliations in terms of position-almost-sameness are a reflection of the correspondence between position and momentum degrees of freedom, also known as *canonical conjugation*, that is one of the characteristic features of the Hamiltonian formulation of classical mechanics.

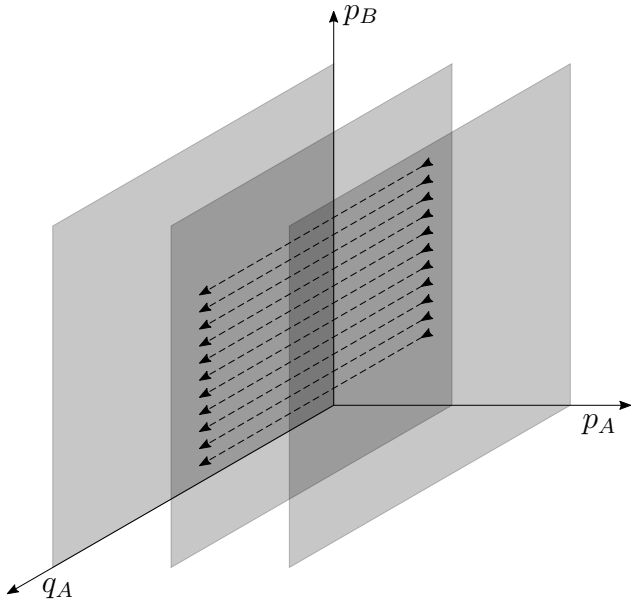


Figure 2.4: The flow lines (dotted) of the Hamiltonian vector field of p_A.

about particle positions, and any two hereditarily momentum-almost-same states agree about the positions of all particles but possibly disagree about particle momenta. Next, let a *position-degenerate subspace* be a subspace of points hereditarily momentum-almost-same to some degenerate point and let a *momentum-degenerate subspace* be a subspace of points hereditarily position-almost-same to some degenerate point. Individualistically speaking, states in position-degenerate subspaces agree about the spatial location at which all particles coincide but can disagree about particle momenta, and states in momentum-degenerate subspaces agree about the momentum value instantiated by all particles but can disagree about particle positions.[11]

The strategy is now to identify individualistic quantification over particles with quantification over position- and momentum-foliations, and quantification over position/momentum values with quantification over position-/momentum-degenerate subspaces. Here is a first pass at this idea: at t, some particle has spatial location q iff there is a primitive state x such that x is actual at t, q is a position-degenerate subspace and there is a position foliation Q

[11] When the symplectic manifold that corresponds to state space is a cotangent bundle, position-degenerate subspaces correspond to certain specific *fibres* in this bundle and momentum-degenerate subspaces to certain specific *sections*.

such that Qxz for every point z in q.

This is not quite right. Consider the claim that, at t, some particle has spatial location q *and* momentum p. The first pass suggests that this claim is equivalent to the proposition that

1. there is a primitive state x such that x is actual at t *and*
2. q is a position-degenerate subspace and p is a momentum-degenerate subspace *and*
3. for some position foliation Q and some momentum foliation P: Qxy for every y in q and Pxz for every z in p.

But this is incorrect. The individualistic proposition attributes position and momentum values to the same particle. However, there is no corresponding structure in its proposed fundamental counterpart: something which serves as the ground for individualistic attributions of positions and momenta, an entity in virtue of which the fundamental structure of the classical symmetry-fundamentalist world is rich enough to give rise to non-fundamental individualistic profiles. What is missing, in individualistic terms, is a fundamental notion which ensures that position- and momentum-foliations—the fundamental surrogates for position and momentum values—belong to the same particle. The state-space-first metaphysician needs to make an additional posit, and one straightforward option is the following relation:

> **4. Particle-Sameness.** A dyadic relation which relates a position- and momentum-foliation iff they belong to the same particle.

The role played by this relation is analogous to imposing a convention that details which phase-space coordinates belong to which particle. For an N-particle system, the standard convention is that the first N triples of phase-space coordinates of a given phase-space point are the position coordinates and the last N triples are the momentum coordinates, so that coordinates $1, 2, 3$ and $N+1, N+2, N+3$ are the position and momentum coordinates, respectively, of the same particle. The relation of particle-sameness is a fundamental posit that captures just this kind of structure at the fundamental level.

Let me stress again that neither the label nor the individualistic characterization of particle-sameness should be read as a metaphysical definition of this relation, or as implying that a state-space-first metaphysics must ultimately reintroduce particles at the fundamental level. Just like almost-sameness and degeneracy, particle-sameness is fundamental and so cannot be defined in more fundamental terms. Nonetheless, this fact does not prevent us from using individualistic language to pin down these fundamental notions—as long as we ultimately succeed in vindicating such individualistic characterizations.

Indeed, this is precisely what particle-sameness allows us to achieve. The result is the following: at t, some particle is at position q with momentum p iff

1. there is a primitive state x such that x is actual at t *and*

2. q is a position-degenerate subspace and p is a momentum-degenerate subspace *and*

3. for some position foliation Q and some momentum foliation P such that P and Q belong to the same particle: Qxy for every z in q and Pxz for every z in p.

Let me illustrate this proposal for a world containing two particles in one spatial dimension. (For ease of exposition, I will label these particles as A and B. Nothing of metaphysical significance hangs on these labels.) The state space of this world is four-dimensional. Figure 2.5 shows a projection of this space onto a three-dimensional subspace in which the position of particle B is fixed. In other words, this diagram depicts some specific leaf of a position-foliation. There are now three things to explain.

First of all, the degenerate states form a two-dimensional subspace. Since one dimension is suppressed in figure 2.5, the degenerate states form a one-dimensional subspace. Since no degenerate state in this diagram can differ about the fixed position of particle B, this line must intersect the q_A-axis at a right angle. Moreover, states in this subspace agree in the momenta of both particles, so this line is at a 45° angle to both the p_A and the p_B axes. If a is the state at which q_A matches q_B, then this subspace is given by the dashed line.

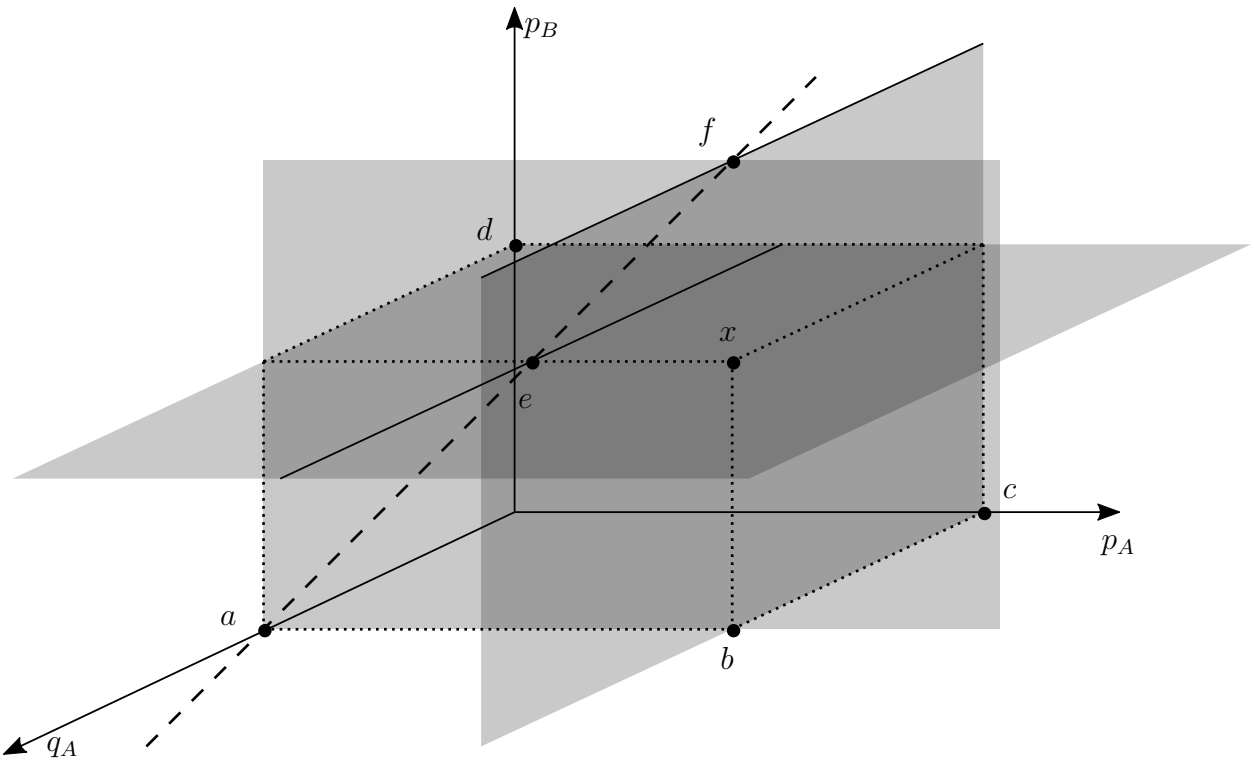


Figure 2.5: Position- and momentum-degenerate subspaces as well as position- and momentum-foliations of a two-particle world in one spatial dimension.

Second of all, position- and momentum-degenerate subspaces are two-dimensional subspaces of state space, while the leaves of position- and momentum-foliations (or position- and momentum-leaves, for short) are three-dimensional subspaces. Position- and momentum-leaves are projected down to two-dimensional surfaces shown in this diagram: the surfaces spanned by *dex* and *bcx* are leaves of momentum-foliations, respectively, and *abx* is a leaf of a position foliation. Since states in this diagram can disagree at most in the position of A, momentum-degenerate subspaces are one-dimensional, such as the solid lines intersecting states f and e. By contrast, position-degenerate subspaces in this diagram are two-dimensional surfaces that (due to the suppression of one dimension) coincide with leaves of position foliations. An example is the surface spanned by *abx*.

Finally, we can verify the above account of individualistic propositions. At state x, the position of particle A is the position-degenerate subspace that coincides with *abx* and its momentum is the momentum-degenerate subspace through f: the former coincides with the

position-leaf spanned by abx and has the same q_A-coordinate as x, the latter is completely within the momentum-leaf spanned by bcx and has the same p_A-coordinate as x. Similarly, the momentum of particle B is the momentum-degenerate subspace intersecting e, which lies completely within the momentum-leaf spanned by dex and has the same p_B-coordinate as x. The relation of particle-sameness is not pictured in this diagram; think of it as fixing the particle labels of the coordinate axes.

The proposal just illustrated is sufficient for specifying the fundamental propositions that entail the individualistic propositions. Complete individualistic profiles are conjunctions of existentially quantified claims of the form 'there is some particle at position q with momentum p'. As we just saw, each of those conjuncts is equivalent to a fundamental proposition about primitive states and the properties of and relations among states. We have therefore met the key challenge facing the state-space-first metaphysician: to specify fundamental propositions (about primitive states and their relations) such that every complete individualistic profile is entailed by some such proposition. On the resulting view, the complete individualistic profile of the world at some instant of time is equivalent to a proposition that details which primitive state is actualized at that time as well as the pattern in the primitive physical structure on state space.

The following table contains a list of the ontological primitives of state-space-first metaphysics in its state-space substantivalist incarnation:

State-Space Substantivalism	
Ontological Primitive	*Type*
Primitive states.	A collection of first-order individuals.
Instants of time.	A collection of first-order individuals.
Actuality.	A single dyadic relation that relates primitive states and instants of time.
The properties and relations that confer symplectic structure on primitive states.	A collection of properties of and relations among primitive states.
The properties and relations that confer temporal structure on instants of time.	A collection of properties of and relations among instants of time.
Position- and momentum-almost-sameness.	Two dyadic relations among primitive states.
Degeneracy.	A monadic property of primitive states.
Particle-sameness.	A dyadic second-order relation that relates dyadic relations among primitive states.[12]

State-space-first proposals like state-space substantivalism give symmetry fundamentalism a precise metaphysical underpinning. As mentioned earlier, state-space symmetries—phase space flows that preserve the symplectic form—correspond to physical entities, the almost-sameness relations. State-space-first proposals of this kind therefore make good on the idea that mathematical symmetries correspond to fundamental aspects of physical reality, whereas things like particles are ontologically derivative of these aspects.

Let me briefly dispel a worry that some readers might have about the kind of ontology presented in this chapter: since states seem to have modal characteristics, the view that state space is fundamental looks appears to imply a kind of modal realism, according to which

[12]In the language of relational type theory, a first-order property is an entity of type $\langle e \rangle$; a first-order relation is an entity of type $\langle e, e \rangle$; and a second-order relation among dyadic first-order relations is an entity of type $\langle \langle e, e \rangle, \langle e, e \rangle \rangle$; where e is the type of individuals, $\langle\rangle$ the type of propositions and for any types $\tau_1, ..., \tau_n$, $n \geq 0$, $\langle \tau_1, ..., \tau_n \rangle$ is a type; nothing else is a type. See (Dorr, 2016, pp. 49-52) for details.

non-actualized, *merely possible* states are every bit as real as the actualized state.

Contrary to appearances, state-space substantivalism does not have this implication. As already noted, *individualistic* states certainly do have a modal flavor: for every proposition about the positions and momenta of particles statable in the individualistic theory, complete individualistic profiles entail either it or its negation. Intuitively, an individualistic state can be thought of as representing an instantaneous way the world could be—much like a possible world. But these modal characteristics of individualistic states are not inherited by their counterparts in the state-space-first ontology, the primitive states. As I argued earlier, the fact that the world is in some primitive state at some instant of time by itself does not fix the complete individualistic profile of the world at that instant; for that, we also need to know the relationships of the primitive state to the other primitive states. According to state-space substantivalism, the space of primitive states is simply the fundamental space in which the history of the world unfolds.

To be sure: the fundamental state space does obey the constraint, familiar from the individualistic setting, that for any instant of time, the world can be in no more and no less than one state at that instant. But whereas in the individualistic theory, this constraint follows from the fact that individualistic states are maximally specific and exhaustive, the state-space fundamentalist theory must add this constraint as a basic law: there is nothing about the nature of primitive states that guarantees mutual exclusivity. Any kind of modal status that this constraint may be thought to confer on primitive states has its source in this basic law rather than in some putative intrinsic modal property of primitive states.

In this chapter, I developed an ontology for classical particle mechanics that gives symmetry fundamentalism a precise metaphysical underpinning—the view that symmetries correspond to fundamental aspects of physical reality, whereas the entities we might ordinarily have thought of as the fundamental building blocks of the world, such as particles or fields, are ontologically derivative of those aspects. This case study serves as a guide for articulating symmetry-fundamentalist ontological proposals in the quantum setting. The next chapter

will articulate such an account for non-relativistic quantum mechanics.

Given the absence of coherent ontological alternatives, individualism has so far been the default option for the ontology of classical particle mechanics. One result of this chapter is that this is no longer the case: the foregrounding of symmetries in metaphysical theorizing led us to state-space-first views as a range of coherent ontological alternatives to individualism. It may therefore be worth reopening the question whether individualism should be regarded as the superior metaphysics for classical particle mechanics. At the very least, a thorough and systematic evaluation of the strengths and weaknesses of state-space-first proposals *vis-à-vis* individualism might improve our grasp of why the latter is preferable as an ontology for classical particle mechanics. This is a task for another day.

Chapter 3

Symmetry Fundamentalism in Quantum Mechanics

The previous chapter articulated an instance of symmetry fundamentalism in the context of classical particle mechanics. The goal of this chapter is to do the same for non-relativistic quantum mechanics. (I'll say more about what I mean by 'non-relativistic quantum mechanics' in a moment.)

I said earlier that the symmetry fundamentalist must give an account of how every proposition about the world's material constituents statable in the materialistic theory is entailed by some fundamental proposition about the ontological counterparts of symmetries. In the classical setting, the materialistic theory is individualism, the theory of point particles located in physical three-space. As I briefly noted at the beginning of the previous chapter, I will assume that, in quantum mechanics the relevant family of materialistic theories consists in those that posit a fundamental field corresponding to the wavefunction. According to these theories, this field is either the only fundamental material constituent of the world, or it is *among* a sparse collection of fundamental material constituents.

But the challenge of explaining how every relevant proposition about this fundamental field is entailed by some proposition about the counterparts of symmetries is distinct from

the challenge of giving an account of how the higher-level objects such as tables, chairs, and measurement devices are constituted out of the fundamental field (plus whatever other ontological items are posited). This challenge, also known as the quantum measurement problem, is beyond the scope of the present investigation: as much as it is besides the point of the previous chapter to explain how tables and chairs can be built from point particles, it is besides the point of this chapter to explain how tables and chairs can be built from the wavefunction (plus whatever other ontological items might be posited by the theory in question). Aside from a few brief remarks at the end of this chapter, this dissertation will not have anything new to say about the specific problem of how to recover the 'manifest image' from the fundamental quantum-mechanical realm.

The project of developing symmetry-fundamentalist ontology for quantum theory is overdue. While the paradigm shift towards symmetries has had an enormous influence on recent fundamental physics, it has had virtually no impact on metaphysical theorizing outside of ontic structuralist circles.[1] Most of the literature on the metaphysics of modern physics has focussed on the ontology of non-relativistic quantum theory—a framework that dates back to the 1920s. Moreover, as I'll explain near the end of this chapter, none of the major views that have resulted from this endeavor—primitive ontology, wavefunction monism, spacetime-state realism—are consistent with symmetry fundamentalism. By contrast, the ontology I develop not only makes good on the idea that symmetries are fundamental but also points us towards a metaphysics for cutting-edge physical theories.

I proceed as follows. In section (3.1) I develop *ray substantivalism*, an ontology for non-relativistic quantum mechanics that vindicates symmetry fundamentalism. I then taxonomize possible versions of ray substantivalism, including primitive-ontology-style versions, in section (3.4). Ray substantivalism amounts to a novel ontology for non-relativistic quantum mechanics. In section (3.5), I contrast this ontology with major extant ontological views about quantum mechanics: wavefunction monism, primitive ontology, and spacetime state realism.

[1] I briefly comment on ontic structuralism in the final chapter of this dissertation.

3.1 Symmetry-Fundamentalist Ontology

In this section, I develop a symmetry-fundamentalist ontology for non-relativistic N-particle quantum mechanics. The mathematical framework I take as the starting point consists of two things: first, a Hilbert space whose elements are used to characterize the quantum states of the system. Second, a range of Hermitian operators that encode the physical structure among quantum states; more on this in a moment. Call this formalism the *Hilbert space framework*.[2] Some version of it is part of the formalism of virtually every existing theory of non-relativistic quantum mechanics, including Bohmian mechanics, Everettian quantum mechanics, and dynamical-collapse theories.

In this setting, state-space symmetries are functions from Hilbert space to itself—also known as *Hilbert-space automorphisms*—that preserve the inner product, i.e. that preserve the characteristic structure of Hilbert space. To vindicate symmetry fundamentalism, we need to articulate an ontology that makes transparent what physical entities correspond to these automorphisms. This is what I turn to now.

A Hilbert-space automorphism can be thought of as a relation among Hilbert space vectors: the relation that obtains between vectors x and y iff the automorphism maps x to y. A natural way to implement fundamentalism about state-space symmetries is therefore to posit fundamental relations that correspond to these automorphisms.

This suggests that the relata of those relations should be regarded as fundamental entities as well. On a straightforward understanding of fundamentality, an entity is fundamental *if and because* it figures in some fundamental proposition.[3] Thus, if the physical counterpart of a state-space symmetry is a fundamental relation, then this is so because it occurs in some fundamental proposition. But the sorts of propositions we express in terms of state-

[2] A further step involves the specification of a differential equation that determines which curves through Hilbert space correspond to the dynamically allowed histories of the quantum state: for example, the Schrödinger equation or a dynamical-collapse equation. Since the symmetries at issue in this paper are state-space symmetries rather than dynamical symmetries, this aspect of the formalism is not directly relevant to the project in this chapter.

[3] (Schaffer, 2009; Audi, 2012; Dasgupta, 2014).

space symmetries also mention the *relata* of those relations: whatever physical entities correspond to Hilbert space vectors. On this line of thought, the fundamentality of the relations corresponding to state-space symmetries implies the existence of a fundamental space whose points are the relata of those relations.

It is natural to think of this fundamental space as *state space*, the space of quantum states of the world. In the context of non-relativistic N-particle quantum mechanics, this raises an important choice point. In physics as well as in the foundations of physics, it is standard to regard quantum states as corresponding to Hilbert space rays rather than to Hilbert space vectors. On this view, state space corresponds not to linear Hilbert space but to *projective* Hilbert space.

Notwithstanding this standard assumption, physicists and philosophers alike routinely appeal to linear Hilbert space structure in their theoretical inferences; indeed, reformulating the entire technical machinery of quantum mechanics into the language of projective Hilbert spaces is a challenging task in its own right.[4] Articulating an ontology that is fundamentalist about projective Hilbert space therefore faces obstacles that go beyond the usual difficulties of metaphysical theorizing about some theory of mathematical physics: it needs to engage with (and perhaps even develop new aspects of) a mathematical machinery that is non-standard in both physics and philosophy. In light of these considerations, some may be moved to more seriously entertain a seemingly naïve ontology that breaks with the consensus and posits states that stand in one-to-one correspondence to Hilbert space vectors.

Although the project of developing a state-space fundamentalist ontology for projective Hilbert space is beyond the scope of this dissertation, the view I propose strikes a compromise between this ambitious project and naïve view. According to the ontology I call *ray substantivalism*, there is a fundamental space that has the structure of linear Hilbert space. But physical states are identified not with points in this space but instead with certain mereological constructions out of these points—entities that are in one-to-one correspondence with

[4]This is evidenced by the approach in (Ashtekar and Schilling, 1999).

rays. To develop this view, I will first give an explicit statement of the more simple-minded ontology, a view I call *vector substantivalism*. This will serve as the basis for the subsequent articulation of ray substantivalism.

It is worth stressing right at the start that vector and ray substantivalism are importantly different from proposals that go under the banner of *wavefunction realism*. According to these proposals, the main role of Hilbert space is to characterize the physical fields on some underlying fundamental space, fields that are mathematically charaterized in terms of the *wavefunction*. Within the Hilbert space framework, the wavefunction of an abstract Hilbert space 'ket' $|\varphi\rangle$ is defined as $\varphi(q_1, ..., q_N) = \langle q_1, ..., q_N|\varphi\rangle$, where $|q_1, ..., q_N\rangle$ is a simultaneous eigenvector of the position operators and N is the number of particles.[5] Vector and ray substantivalists hold a different view of the role of Hilbert space: instead of being a way of characterizing physical fields on some underlying fundamental space, Hilbert space corresponds to a fundamental space in its own right: the space whose physical constituents are in one-to-one correspondence with abstract Hilbert space vectors.

3.2 Vector Substantivalism

With those clarifications out of the way, we can move on to the first step: a precise statement of vector substantivalism. This ontology consists of two ranges of first-order individuals: first, a collection of entities that stand in one-to-one correspondence with Hilbert-space vectors. I'll henceforth refer to these entities as *Points*. The second range of first-order individuals are instants of time.

According to vector substantivalism, Points are quantum states; that is, they are the

[5]Strictly speaking, eigenvectors of unbounded Hermitian operators on $\mathcal{H}$ (such as those associated with the position and momentum degrees of freedom) are not elements of $\mathcal{H}$ but of the *rigged* Hilbert space associated with $\mathcal{H}$; i.e the triple consisting of $\mathcal{H}$, a dense subspace $\Phi \subset \mathcal{H}$, and its dual Φ^* such that $\Phi \subset \mathcal{H} \subset \Phi^*$. In this more rigorous setting, position and momentum eigenstates are 'bras', complex linear functionals that live in Φ^* and that combine with 'kets' to yield a complex number. Here, I will follow the vast majority of discussions in the philosophy of quantum mechanics and ignore these technicalities. However, none of these technical complexities seem to be insurmountable obstacles: it should be possible in principle to specify ontological primitives that capture the structure of a rigged Hilbert space. This is a project I will not pursue here.

sorts of entities that can be 'actualized' at various instant of time. This is captured by the following posit:

> **Point actuality.** A dyadic relation. Relates a Point x and an instant of time t iff x is actualized at t.

For Points to behave like states, the vector substantivalists posits the following basic law familiar from the previous chapter: necessarily, for every instant of time there is some Point (and no more than one Point) that is actualized at it.[6,7]

Over and above Points, times, and actuality, the vector substantivalist requires posits that confer on Points the structure of a Hilbert space. A fully nominalistic account of this structure is beyond the scope of this chapter.[8] Here I will assume that, as far as state-space structure on Points is concerned, we can help ourselves to ontological primitives that take complex numbers as arguments. Together with standard Hilbert space axioms, the following primitives confer Hilbert space structure on states:

> **1. Sum.** A triadic relation. Relates Points x, y, and z iff x is the sum of y and z. Write $x = y + z$.
>
> **2. Multiplication.** A triadic relation. Relates Points x and y and a complex number c iff x is a multiple by c of y. Write $x = cy$.

[6]Since ray substantivalism quantifies over times, it presumes something like a B-theoretic account of time. An A-theoretic variant of ray substantivalist can be obtained by replacing quantification over instants of time with propositional tense operators such as 'it will always be the case that...', 'it will at some time be the case that...' etc, and replacing the Point actuality relation with a monadic property '...is actual'.

[7]Ray substantivalism is not the only way to implement fundamentalism about the physical counterparts of Hilbert space vectors. I already remarked in section 2.1 that there are various options for how to conceive of primitive states in classical mechanics; similar options exist in the quantum setting as well. The first alternative is to posit a fundamental first-order individual, 'the world', and to think of Points as *relations* such that necessarily, for every instant of time, there is some such relation that obtains between the world and that instant. The A-theoretic version of this view regards Points as fundamental monadic properties such that, necessarily, the world always instantiates some Point. A second alternative is to regard Points as monadic properties of instants of time such that, necessarily, every instant of time instantiates some such property. The A-theoretic version can be thought of as a kind of *propositionalism*: the view that Points are propositions or states of affairs such that, necessarily, there is always some such proposition that is true or some such state of affairs that obtains. This choice point raises deep foundational issues that are beyond the scope of this chapter; cf. (Bacon, 2019) for some pertinent discussions.

[8]See (Chen, 2018) for recent work on nominalizing quantum theory.

3. Inner Product. A triadic relation. Relates Points x, y and a complex number c iff c is the inner product of x and y. Write $c = (x, y)$.

The main challenge for the vector substantivalist is to show that, for every proposition that details the pattern in the physical entity corresponding to wavefunctions, there is some fundamental proposition about Points that entails it. Indeed, we will see that something even stronger is true: every proposition that details the pattern in the physical field corresponding to the wavefunction is *equivalent* to some fundamental proposition about Points.

It should be clear that the ontological primitives introduced so far are insufficient for this task. Different spaces of Points can agree about the pattern in the properties and relations that confer Hilbert space structure while disagreeing in the kinds of ontologically derivative constituents of the world. A Hilbert space of appropriate dimensionality can be used both to describe a world that contains physical systems with spin and a world that contains physical systems without spin. The vector substantivalist therefore requires additional primitives.

I will introduce the relevant primitives in terms of their mathematical characterizations; but it as I emphasized in the previous chapter, these characterizations are not metaphysical definitions. It should go without saying that ontological primitives cannot be defined in more fundamental terms, let alone in terms of their mathematical description. As is common in metaphysical theorizing, novel fundamental posits must be justified by exhibiting their metaphysical role in an account that highlights their explanatory and inferential connections to more familiar notions. This is precisely my goal here: it will shortly become clear that the primitives figure in an account according to which every relevant proposition about the counterparts of wavefunctions is entailed by some proposition about these primitives.

To introduce these posits, I need to revisit a remark I made earlier: that Hermitian operators on Hilbert space are used to encode the physical structure among quantum states. This idea is most transparent and familiar from the wavefunction realist perspective. We already saw how the wavefunction is defined in terms of the eigenvectors of the position operators: the wavefunction of any ket $|\varphi\rangle$ is $\varphi(q_1, ..., q_N) = \langle q_1, ..., q_N|\varphi\rangle$, where $|q_1, ..., q_N\rangle$ is

a simultaneous eigenvector of the position operators. This procedure doesn't just define some particular wavefunction; it defines *all possible* wavefunctions.[9] In other words, we obtain a range of wavefunctions that characterize the possible instantaneous arrangements of the physical field over the fundamental space, arrangements that possibly differ in the property that the field assigns to a given location in this space. In this sense, Hermitian operators designated as 'position operators' encode the physical structure of the fundamental field posited by the wavefunction realist.

In the vector substantivalist setting, the thought is quite similar. Every Point can be mathematically characterized in terms of a linear combination of eigenvectors of the relevant Hermitian operators. For example, a ket $|\varphi\rangle$ can be specified in terms of position eigenvectors as

$$|\varphi\rangle = \int_{-\infty}^{\infty} |q_1, ..., q_N\rangle\langle q_1, ..., q_N|\varphi\rangle dq_1 ... dq_N. \tag{3.1}$$

Every difference between distinct kets therefore corresponds to a difference in some complex coefficients that figure in the expansions of those kets in terms of the simultaneous position eigenvectors. This illustrates how Hermitian operators encode *physical relations* among distinct Points: respects in which Points possibly differ from each other.

One possible strategy for the vector substantivalist is therefore to posit relational primitives that capture the structure of Hermitian operators on Hilbert space. Here, I follow a different path: one that leads to an instance of symmetry fundamentalism. For ease of exposition, let me introduce some terminology. Say that each of the N real numbers that figure in simultaneous position eigenvectors such as $|q_1, ..., q_N\rangle$ are values of the *position variables*. Correspondingly, for any $|q_1, ..., q_N\rangle$ and $|q'_1, ..., q_N\rangle$ such that $q_1 \neq q'_1$, I will say that the two vectors *differ in the first position variable*. This extends to arbitrary linear combinations of

[9] For systems with spin, the Hermitian operators corresponding to the spin degrees of freedom characterize the *kind* of property that the fundamental field assigns to every point: namely, properties mathematically specified in terms of spinors rather than in terms of complex numbers.

eigenvectors. For example, the following vectors differ only in the first position variable:

$$|\varphi\rangle = \int_{-\infty}^{\infty} |q_1, ..., q_N\rangle\langle q_1, ..., q_N|\varphi\rangle dq_1...dq_N \tag{3.2}$$

$$|\phi\rangle = \int_{-\infty}^{\infty} |q_1 + a, ..., q_N\rangle\langle q_1, ..., q_N|\varphi\rangle dq_1...dq_N. \tag{3.3}$$

Similarly for momentum. So, there is a position/momentum variable for every position/momentum operator.

We now take note of the following fact: specifying a range of Hermitian operators on Hilbert space is mathematically equivalent to specifying a range of state-space symmetries. More specifically, due to a result known as Stone's theorem, there is a one-to-one correspondence between Hermitian operators and one-parameter unitary transformations.[10] The mathematical details do not matter for our purposes; what's important is that, if $\hat{p}$ is a momentum operator,

$$U(a) = e^{ia\hat{p}} \tag{3.4}$$

is the unitary operator that implements a shift by amount a in the corresponding *position* variable. Similarly, if $\hat{q}$ is a position operator,

$$V(k) = e^{ik\hat{q}} \tag{3.5}$$

is the unitary operator that implements a shift by amount k in the corresponding *momentum* variable. For example, the unitary operator $U_1(a)$ that acts on the simultaneous position eigenvector $|q_1, q_2, ..., q_N\rangle$ as $U_1(a)|q_1, q_2, ..., q_N\rangle = |q_1 + a, q_2, ..., q_N\rangle$ implements a shift in the first position variable. The action of shift operators on eigenvectors determines their action on arbitrary linear combinations of eigenvectors. For example, it is immediate from equations

[10](Stone, 1932).

(2) and (3) that

$$|\phi\rangle = U_1(a)|\varphi\rangle. \tag{3.6}$$

The unitary position and momentum shift operators suggest two ontological primitives that capture the fundamental relational structure on Points. The idea is that we posit one fundamental relation that obtains between two Points when they are related by a shift in some position variable, and one fundamental relation that obtains between two Points when they are related by a shift in some momentum variable:

> **4. Position-Almost-Sameness.** A dyadic relation. Relates two Points just in case they agree about all position variables except for a shift in one position variable.
>
> **5. Momentum-Almost-Sameness.** A dyadic relation. Relates two Points just in case they agree about all momentum variables except for a shift in one momentum variable.

(Note that these are relations are distinct from the relations by the same label introduced in the previous chapter. I will continue to use these labels throughout this chapter for ease of exposition and to highlight the significant similarities between the classical and the quantum case.)

I will show now that, for systems without spin, these two relations (together with the relations that confer Hilbert space structure on Points) are sufficient for specifying a range of fundamental propositions such that, for every proposition about the ontological counterparts of wavefunctions, one of these fundamental propositions entails it.

My strategy takes inspiration from the above-mentioned mathematical definition of wavefunctions in terms of kets: $\phi(q_1, ..., q_N) = \langle q_1, ..., q_N|\phi\rangle$. The main idea is that, for any fixed ket, this expression can be thought of as a relation between simultaneous eigenvectors and the complex numbers. This suggests that the ontological counterpart of the wavefunction

is some specific pattern in the inner product relation. To implement this idea, we need to define two pieces of derivative ontology: first, a property instantiated by a Point just in case its mathematical representative is a simultaneous eigenvector; second, physical entities that are surrogates for points of configuration space and points of momentum space, respectively.

I start with the ontological counterpart of a simultaneous eigenvector. Let a *position-almost-same subspace* be a space of Points such that any two Points in it are position-almost-same; and *momentum-almost-subspace* is a space such that any two Points in it are momentum-almost-same. By an argument analogous to the classical case, we can show that Points in a given position- or momentum-almost-same subspace differ about the same position or momentum variables. Let y, z be Points that are position-almost same to some Point x, and suppose that the differences between x, y and the differences between x, z concern *different* position variables. It follows that y, z differ about more than one position variable and so at most one of y and z can be in the same position-almost-same subspace as x. (Similarly for momentum-almost-same subspaces.)

Next, we observe that, if $|p\rangle$ is an eigenvector of a momentum operator $\hat{p}$ then it is also an eigenvector of the unitary operator $U(a)$ that implements shifts in the corresponding position variable:

$$U(a)|p\rangle = e^{iap}|p\rangle \tag{3.7}$$

A similar fact holds for position eigenvectors: if $|q\rangle$ is an eigenvector of a position operator $\hat{q}$ then it is also an eigenvector of the unitary operator $V(k)$ that implements shifts in the corresponding momentum variable:

$$V(k)|q\rangle = e^{ikq}|q\rangle \tag{3.8}$$

In other words: kets that are eigenvectors of some momentum operator are such that the spatial shift operator maps them to complex scalar multiples of themselves; and momentum

eigenvectors are the only kets with this property. Similarly, kets that are eigenvectors of some position operator are such that the momentum shift operator maps them to complex scalar multiples of themselves; and position eigenvectors are the only kets with this property. We can use this fact to define the following physical properties that correspond to the notion of an eigenvector:

> **A. Position-Eigenpoint.** A monadic property. Instantiated by a Point x iff for some momentum-almost-same subspace containing x, every Point in that subspace is a multiple of x.
>
> **B. Momentum-Eigenpoint.** A monadic property. Instantiated by a Point x iff for some position-almost-same subspace containing x, every Point in that subspace is a multiple of x.

The properties of being a position- or momentum-eigenpoint are instantiated by Points whose mathematical surrogates are eigenvectors of at least one position or momentum operator: for example,

$$|q_1, \varphi_2, ..., \varphi_N\rangle = \int_{-\infty}^{\infty} |q_1, q_2, ..., q_N\rangle \varphi_2(q_2)...\varphi_N(q_N) dq_2...dq_N \tag{3.9}$$

is an eigenvector in the first position variable but not necessarily an eigenvector in other position variables. But to define the intrinsic counterparts of position and momentum wavefunctions, we also need properties instantiated by Points when their mathematical surrogates are *simultaneous* eigenvectors of all position or all momentum operators. These can be defined as follows:

> **C. Simultaneous Position-Eigenpoint.** A monadic property. Instantiated by a Point x iff every Point momentum-almost-same to x is a multiple of x.
>
> **D. Simultaneous Momentum-Eigenpoint.** A monadic property. Instantiated by a Point x iff every Point position-almost-same to x is a multiple of x.

The second task is to define two ranges of entities that act as surrogates for points of configuration space and points of momentum space, respectively. It may be tempting to think that simultaneous position- and momentum-eigenpoints are good candidates for this role. This temptation should be resisted: there are *too many* simultaneous position and momentum eigenpoints. This is most transparent at the level of mathematical description. Given some fixed coordinate system on configuration space, there is a one-to-one correspondence between points of configuration space and the eigenvalues of the simultaneous position-eigenvectors of position operators along the axes of this coordinate system. But there are many more distinct simultaneous position eigenvectors than there are distinct position eigenvalues: for any complex number c, $|q_1, q_2, ..., q_N\rangle$ and $c|q_1, q_2, ..., q_N\rangle$ are not just both simultaneous position eigenvectors; they are eigenvectors with identical position eigenvalues. This problem can be sidestepped using tools from mereology:

> **E. Configuration space point.** A monadic property. Instantiated by q iff q is a fusion of Points and every atomic part of q is a simultaneous position eigenpoint and any two atomic parts of q are related by a complex scalar multiple.
>
> **F. Momentum space point.** A monadic property. Instantiated by p iff p is a fusion of Points and every atomic part of p is a simultaneous momentum eigenpoint and any two atomic parts of p are related by a complex scalar multiple.

This gives the right result: for every position/momentum eigenvalue, there is a unique configuration/momentum space point: the point such that any two of its atomic parts have this eigenvalue.

We are almost ready to define the ontological counterparts of wavefunctions. The last observation we need to make is that wavefunctions are defined in terms of *normalized* simultaneous eigenvectors, rather than in terms of their complex scalar multiples. For example, we define the wavefunction of $|\phi\rangle$ in terms of vectors like $|q_1, q_2, ..., q_N\rangle$ rather than in terms of $c|q_1, q_2, ..., q_N\rangle$ where $c \neq 1$.

It is important that we take note of this fact. Recall: our strategy is to identify wavefunc-

tions with certain specific patterns in the inner product relation. But whereas wavefunctions assign a complex number to every point of configuration or momentum space, the inner product takes Points as arguments, rather than the fusions of Points that act as surrogates for configuration and momentum space points. This means that, for any given Point x and configuration space point q, the atomic parts of q differ in the complex number assigned to them by the inner product. This raises the question as to which complex number the putative ontological counterpart of a wavefunction should assign to q. The answer is suggested by the mathematical observation in the previous paragraph: it is the number that bears the inner product relation to every *normalized* atomic part of q.

More precisely, let a Point x be *normalized* iff there is some complex number c of unit modulus such that $c = (x, x)$. The ontological counterparts of wavefunctions are then defined as follows:

> **G. Configuration-Space Field of x.** A dyadic relation. Relates a complex number c and a point of configuration space q just in case, for every normalized atomic part y of q, $c = (y, x)$.
>
> **H. Momentum-Space Field of x.** A dyadic relation. Relates a complex number c and a point of momentum space p just in case, for every normalized atomic part y of p, $c = (y, x)$.

According to vector substantivalism, the non-fundamental proposition we express when we specify the wavefunction of some Point x is the proposition that details the pattern in the configuration-space field of x: that is, the proposition that specifies, for every point of configuration space q, the complex number that is the value of the configuration-space field of x at q. (Similarly for momentum.) Since the configuration-space field of x is defined in terms of the primitives (1-5), it follows that every non-fundamental proposition about the ontological counterparts of wavefunctions statable in the Hilbert space framework is equivalent to some fundamental proposition.

Our discussion has two results. First, vector substantivalism entails that there are physical

counterparts of state-space symmetries: the almost-sameness relations. Second, vector substantivalism satisfies the desideratum that, if the ontological counterparts of wavefunctions are ontologically derivative of almost-sameness relations, then every relevant proposition about the former is entailed by some proposition about the latter. Vector substantivalism can therefore be regarded as an instance of symmetry fundamentalism.

3.3 Ray substantivalism

Vector substantivalism assumes that Points are quantum states and is therefore at odds with the view, widespread among physicists and philosophers, that quantum states correspond to rays in Hilbert space rather than to Hilbert space vectors. In this section, I sketch *ray substantivalism*, an ontology that respects this consensus. This view is not an attempt at the more ambitious project of fundamentalism about projective Hilbert space. Instead, it is the view that quantum states are identified with certain mereological fusions of Points.

Ray substantivalism is a fairly straightforward modification of vector substantivalism. The fundamental first-order posits are identical: they consist in Points and instants of time. Moreover, the ray substantivalist follows the vector substantivalist in positing Hilbert space structure (1-3) among Points as well as the relations of position- and momentum-almost-sameness (4,5). The central difference between the two views is the identification of the sorts of entities that can be actualized at various instants of time; that is, the entities that are identified with quantum states. Let a fusion s of Points be a *State* iff for any two atomic parts x, y of s, x is a multiple of y. According to the ray substantivalist, actuality is a fundamental relation between instants of time and States:

> **State actuality.** A dyadic relation. Relates a State s and an instant of time t iff s is actualized at t.

In addition, the ray-substantivalist version of the basic law governing the actuality relation says that necessarily, for every instant of time there is some State (and no more than one

State) that is actualized at it.

The only other feature of vector substantivalism that the ray substantivalist needs to amend is the definition of configuration- and momentum-space fields. Definitions G and H associate a field with every Point. But for the ray substantivalist, this is too fine-grained: the physical entity that corresponds to wavefunctions should be such that there is one such entity for every State rather than one for every Point.

We can obtain the requisite definitions on the basis of G and H. Roughly, the idea is to define the ray-substantivalist field of a State s as the vector-substantivalist field of the normalized atomic part of s.

> **G'. Normalized Configuration-Space Field of s.** A dyadic relation. Relates a complex number c and a point of configuration space q just in case c is assigned to q by the configuration-space field of the normalized atomic part of s.
>
> **H'. Normalized Momentum-Space Field of s.** A dyadic relation. Relates a complex number c and a point of momentum space p just in case c is assigned to p by the momentum-space field of the normalized atomic part of s.

This completes the outline of ray substantivalism. According to this view, at every instant of time, the propositions about the normalized configuration- and momentum-space fields statable in the Hilbert space framework are entailed by which State is actual at that instant, together with the proposition that details the pattern of properties of and relations among the atomic parts of the actualized State and other Points. Ray substantivalism agrees with vector substantivalism about the fundamental first-order posits, as well as about the fundamental structure on Points. But the two views disagree about the nature of quantum states as well as about the physical counterparts of wavefunctions.

Ray substantivalism counts as an instance of symmetry fundamentalism: according to ray substantivalism, almost-sameness relations are fundamental relations that correspond to state-space symmetries; and every relevant proposition about the physical counterparts of wavefunctions is entailed by some fundamental proposition about these relations.

Some readers might object that ray substantivalism does not seem like a genuine advance over vector substantivalism. The central motivation for identifying quantum states as corresponding to rays rather than to Hilbert space vectors is that no experiment can distinguish between states of a quantum system that differ only by a global phase. Hence, if quantum states are in one-to-one correspondence with Hilbert space vectors, then no possible experiment can determine the quantum state of a physical system. Avoiding this kind of underdetermination seems to be the main motivation of the standard view that quantum states correspond to Hilbert space rays.

The objection might then go as follows: it looks as if ray substantivalism achieves consistency with the standard view only at the cost of reproducing the underdetermination at the level of the atomic parts of States. On ray substantivalism, the actualized State contains a rich structure among its parts, a structure which seems in-principle empirically inaccessible. But this worry is misplaced. According to ray substantivalism, there is no in-principle experimental underdetermination of which State the world is in at any given instant of time. And given that the State of the world at some instant of time is known, the structure among its atomic parts can be inferred simply from the linear Hilbert space structure among Points. Ray substantivalism does not feature the sort of underdetermination that spells trouble for vector substantivalism.

3.4 A Taxonomy of Ray-Substantivalist Views

The ontology I just outlined is not the only version of ray substantivalism. The goal of this section is to provide a taxonomy of the available options.

The range of possible ray-substantivalist views mirrors the varieties of wavefunction realism. There are two important respects in which the latter differ. The first concerns the fundamental space that is the container of the physical field characterized by the wavefunction. *Configuration-space realists* say that the fundamental field lives on configuration space, the

$3N$-dimensional space whose points we can intuitively think of as each fixing a distribution of N classical point particles in three-space. *Multi-field realists*, by contrast, say that the fundamental field assigns a property to every N-collection of points of Euclidean three-space.[11] The second line of disagreement is about whether fundamental reality is exhaustively described by the wavefunction. *Wavefunction monists* say 'yes'. The most influential theories that answer this question in the negative are *primitive-ontology*-style theories. According to them, fundamental reality additionally consists in local beables: physical objects that are confined to a bounded region of the fundamental space.

For example, GRW_0 is a configuration-space-realist version of wavefunction monism according to which the dynamical evolution of the fundamental field is given by a collapse equation.[12] Bohmian mechanics is a primitive-ontology-style view of which there exist both a configuration-space-realist as well as a multi-field-realist version. According to the former, there is a single 'world-particle' in configuration space. According to the latter, there are point-particles in three-space.[13] The two theories agree that the evolution of the particulate content of the world is determined by the fundamental field which itself evolves according to the Schrödinger equation.[14] GRW_f and GRW_m are primitive-ontology-style theories that can be given a configuration-space-realist and a multi-field-realist formulation. Both versions posit physical entities (flashes or mass densities) distributed over the fundamental arena whose temporal evolution is determined by the fundamental field. Finally, there are two wavefunction-monist options for Everettian quantum mechanics (EQM): a configuration-space-realist one and a multi-field-realist one.[15]

[11] Note that the term 'wavefunction realism' is sometimes used exclusively for configuration-space-realist views.

[12] (Albert, 1996; Loewer, 1996; Ney, 2012; North, 2013).

[13] (Belot, 2012; Chen, 2017; Hubert and Romano, 2018); cf. (Forrest, 1988).

[14] Note that there are versions of Bohmian mechanics that are not wavefunction realist. According to one proposal, the wavefunction corresponds to a law-like physical fact; on another, the wavefunction corresponds to a *sui generis* physical entity; (Maudlin, 2013); cf. (Chen, 2019).

[15] Although the official ontology for EQM is spacetime-state realism (more on which below), EQM has often been introduced, developed, and criticized on the assumption of a configuration-space-realist wavefunction-monist ontology; for example, (Brown and Wallace, 2005; Saunders, 2010; Maudlin, 2010; Ney, 2013). To my knowledge, a multi-field wavefunction monist ontology for EQM has not been explicitly developed in the literature.

This taxonomy can be transposed more or less directly to ray-substantivalism. The view I articulated can be regarded as *configuration-space-realist monistic* ray substantivalism. The view has a monistic character: at every instant of time, physical reality is determined entirely by which State is actualized at that time, together with relations between the atomic parts of the actualized State and other Points. Moreover, the view I stated is configuration-space-realist: for every State, the ontologically derivative entity that corresponds to the wavefunction is a field on configuration space. Another possible monistic view is *multi-field-realist* ray substantivalism, in which the main ontologically derivative entity is a multi-field on physical three-space. Both versions of monistic ray substantivalism can be combined with dynamical laws of either the dynamical-collapse or the Schrödinger sort, resulting in versions of GRW_0 and EQM, respectively.

There are also ways to supplement monistic ray substantivalism with a primitive ontology. Here, I will focus on two Bohmian versions of this approach. The first version is what you might call *configuration-space-realist Bohmian ray substantivalism*: the view which (over and above the ray substantivalist primitives) posits a single Bohmian 'world-particle' located in fundamental configuration space. An alternative is *multi-field-realist Bohmian ray substantivalism* which includes the posits of monistic ray substantivalism, plus a range of fundamental point particles located in fundamental physical three-space. On both views, the dynamical evolution of the fundamental particle(s) are determined by the actualized State according to the guidance equation; and the dynamical law governing States is given by the Schrödinger equation. These views differ from monistic ray substantivalism in both their fundamental and derivative ontology: they posit fundamental particles as well as points of the respective fundamental spaces and therefore have no need to lay down definitions of these fundamental spaces in terms of Points.

Although Bohmian versions of ray substantivalism are coherent candidate ontologies, there are three reasons to be skeptical. First, they have costs that they don't share with their wavefunction-realist competitors. The Bohmians we're imagining posit two fundamental

spaces: the space of Points, plus whatever space contains the items of primitive ontology. Wavefunction-realist Bohmians, by contrast, get by with positing just one fundamental space: the arena which contains both the Bohmian particles and the fundamental field. Both kinds of Bohmian ray substantivalism therefore seem less ontologically parsimonious than their wavefunction-realist counterparts.

Second, neither version of Bohmian ray substantivalism is an instance of symmetry fundamentalism. To be sure: these views posit almost-sameness relations among Points and are therefore consistent with the first conjunct of symmetry fundamentalism. But they are inconsistent with the second conjunct. The items of primitive ontology are just that: primitive, metaphysically fundamental, not ontologically derivative of anything else. So, it is not the case that every relevant proposition about Bohmian particles is entailed by some proposition about almost-sameness relations.

Third, there is a systematic relation between the symmetries of the fundamental space that contains the Bohmian particle(s) on the one hand and the relevant state-space symmetries on the other. For example, it is routine to define unitary actions of the rotation group $SO(3)$ on Hilbert space in terms of an action of this group as Euclidean isometries.[16] To be sure: the existence of this kind of mathematical definition does not force the Bohmian ray substantivalist to lay down a corresponding metaphysical definition of the relevant sameness relations on state space. But it does mean that these views face the additional burden of explaining how the almost-sameness relations on the space of Points hang together with the symmetries of the fundamental container of the primitive ontology.

[16]Consider a single-particle system in three-dimensional Euclidean space, characterized in terms of $\mathbb{R}^3$, and let $\hat{\mathbf{x}}$ be the position operator with eigenvalues in $\mathbb{R}^3$: if $\hat{\mathbf{x}}|\vec{x}\rangle = \vec{x}|\vec{x}\rangle$ then $\vec{x} \in \mathbb{R}^3$. The standard action of $SO(3)$ on $\mathbb{R}^3$ is a map $\phi : SO(3) \times \mathbb{R}^3 \to \mathbb{R}^3$, $(R, \vec{x}) \mapsto \phi(R, \vec{x}) = R\vec{x}$ where R are real invertible 3×3 matrices such that $RR^T = \mathbb{1}$ and $detR = 1$. A unitary action $U : SO(3) \times \mathcal{H} \to \mathcal{H}$, $(R, |\psi\rangle) \mapsto U(R)|\psi\rangle$ on the one-particle Hilbert space $\mathcal{H}$ is then defined via $U(R)|\vec{x}\rangle = |R\vec{x}\rangle$.

3.5 Discussion

In this section, I do three things. First, I explain a claim that I made in the introduction: that no extant version of the three major proposals for the ontology of quantum mechanics is consistent with symmetry fundamentalism—neither monistic and primitive-ontology variants of wavefunction realism nor spacetime-state realism. Second, I briefly contrast ray substantivalism with major existing proposals for the ontology of non-relativistic quantum mechanics. Third, I address some objections to views along the lines of ray substantivalism.

First, no extant version of the major proposals is consistent with symmetry fundamentalism. Recall that symmetry fundamentalism is the conjunction of two claims: first, that state-space symmetries correspond to fundamental aspects of physical reality; and second, that the kinds of things we would have ordinarily regarded as the fundamental building blocks of the world, such as particles or fields, are ontologically derative from these fundamental aspects.

Existing proposals are inconsistent with the first conjunct: they do not posit fundamental entities that correspond to state-space symmetries. The most straightforward case is wavefunction realism. Theories of this type posit a fundamental physical field on the relevant fundamental arena, plus (perhaps) the distribution of local beables over this arena. On these views, nothing in the ontology corresponds to state-space symmetries.

Another proposal for the ontology of quantum theory is known as spacetime-state realism, one that differs quite significantly from wavefunction realism. According to this proposal, fundamental physical reality consists in an assignment of a physical property to every region of spacetime: the *quantum state* of that region.[17] Mathematically, each region i is assigned a Hilbert space $\mathcal{H}_i$ and the Hilbert space $\mathcal{H}$ of the world is given by the tensor product of these Hilbert spaces: $\mathcal{H} = \bigotimes_i \mathcal{H}_i$. If the quantum state of the universe at time t is given by the density operator ρ on $\mathcal{H}$, then the quantum state of spacetime region i at t is given by the partial trace $\rho_i = Tr_{i \neq j} \rho$.

There are at least two ways to understand this view. The first is more serious about

[17] (Wallace and Timpson, 2010); cf. (Swanson, 2018).

the idea that quantum states are represented by operators on Hilbert spaces. On this line of thought, for every spacetime region i we posit a fundamental space corresponding to $\mathcal{H}_i$ as well as a fundamental relation on this space corresponding to ρ_i. This version of spacetime-state realism collapses into a form of ray substantivalism: an ontology according to which there is a collection of spaces of Points, together with certain privileged relations among those Points.

The similarity between ray substantivalism and this version of spacetime-state realism becomes even more transparent once we take note of an observation made by Ismael and Schaffer (2016). Understood as fundamental entities, the quantum state assigned to spacetime regions smaller than the entirety of spacetime are ontologically redundant: as the above derivation of ρ_i from ρ suggests, the quantum state of any spacetime region is completely fixed by the state of the world as a whole, but not vice versa. According to what Ismael and Schaffer refer to as 'spacetime-state realism streamlined', the fundamental ontology includes only the quantum states of the world as a whole: exactly what the ray substantivalist refers to as States.

So, ray substantivalism and this version of spacetime-state realism are very similar indeed. The main difference concerns the nature of the physical structure on state space: according to ray substantivalism, this structure consists in almost-sameness relations; according to 'spacetime-state realism streamlined', it consists in the physical properties and relations on state space that correspond to certain Hermitian operators on Hilbert space.[18] Hermitian operators, however, are not themselves state-space symmetries but *generators* of state-space symmetries. So this version of spacetime-state realism is not an instance of symmetry fundamentalism.

This is not the version of spacetime-state realism intended by its main proponents.[19] According to the official version of spacetime-state realism, the fundamental properties

[18](Wallace, 2012, p. 296).

[19]For example, David Wallace has rejected what he refers to as 'Hilbert-space-vector realism' (Wallace, 2017, p. 8) and 'Hilbert space realism' (Wallace, 2013, p. 216).

assigned to spacetime regions correspond to abstract operator algebras. Their technical definition is not important here; what matters is that the operators in these algebras need not be construed as operators acting on a Hilbert space.[20] So this version of spacetime-state realism does not seem to be at risk of collapsing into ray substantivalism. However, it is not an instance of symmetry fundamentalism. Given the absence of Hilbert spaces from this formulation, there is no clear sense in which the operators that encode the relevant physical properties can be thought of as 'acting on' anything, much less as being state-space symmetries.

To be sure: for every Hilbert-space automorphism that counts as a state-space symmetry within the Hilbert-space framework, there is an automorphism of the operator algebra within the spacetime-state realist formalism, and vice versa. However, once the relevant operator algebras associated with each spacetime region have been specified, the symmetries of the algebra are physically redundant, in the sense that the relevant physical structure is already exhaustively specified in terms of the operator algebra alone. Let me illustrate this point by analogy to a fact about the Hilbert-space framework we encountered earlier. Recall that, in this framework, the relevant physical structure on Hilbert space can be specified either in terms of self-adjoint operators or in terms of the unitary transformations they generate. But once we have made the choice (for example) to characterize the physical structure in terms of self-adjoint operators, Stone's theorem implies that nothing is added by also specifying the corresponding unitary transformations: all the physical structure is already fixed by the self-adjoint operators. The situation with regard to this version of spacetime-state realism is similar: once the relevant operator algebras have been specified, there is simply no job to be done by the symmetries of the algebra.[21] The operator-algebra formulation of spacetime-state realism is therefore not properly understood as an instance of symmetry fundamentalism.

The result is that none of the major extant proposals for the ontology of quantum

[20] Cf. (Wallace and Timpson, 2010; Swanson, 2018).

[21] I thank David Baker for pressing me on this point.

theory are instances of symmetry fundamentalism.[22] Insofar as compatibility with symmetry fundamentalism is concerned, monistic ray substantivalism therefore has some claim to be the only game in town. Those who are in the market for a view that provides a precise ontological basis for symmetry-fundamentalism should seriously entertain monistic ray substantivalism as a candidate ontology for non-relativistic quantum theory.

Second, there are aspects of monistic ray substantivalism that might make it less attractive than its competitors. A minimal adequacy condition for an ontology for quantum theory is that it be consistent with what is sometimes referred to as the 'manifest image': the familiar world of medium-sized dry goods, such as tables, chairs, and measurement devices. If an ontology does not satisfy this constraint, it is plainly inconsistent with the empirical evidence. However, some philosophers have suggested that we should impose more demanding constraints on ontological theorizing. An admissible ontology must not just *be consistent with* the manifest image, it must also be *reflective* of certain important aspects of this image. Perhaps the most well-known example of this is the kind of Bohmian approach according to which fundamental ontology ought to reflect the fact that the macroworld seems to consist of local beables. The way in which fundamental ontology is required to reflect this fact is that it should contain local beables as well.[23]

Those who subscribe to constraints of this sort are unlikely to find monistic ray substantivalism attractive. It's not just that the fundamental ontology does not include local beables. Since there is no fundamental physical spacetime, the notion of a spatiotemporally localized entity is not even expressible in fundamental terms.

However, this predicament should be familiar from discussions of configuration-space-realist wavefunction monism, a view originally proposed by David Albert (1996).[24] Both views posit a

[22]Moreover, note that it would not help to put such entities in 'by hand', as it were. Wavefunction realist and spacetime-state realist theories are constructed to be ontologically complete in the sense that every relevant non-fundamental proposition is entailed by some proposition that is just about the fundamental physical field, and perhaps also about the primitive ontology. The physical counterparts of symmetries would therefore be ontological danglers: every relevant non-fundamental proposition could be derived from a proposition that doesn't mention these entities.

[23]See (Allori, 2013) for a succinct expression of this view.

[24]This view is further developed in (Ney, 2012); cf. (North, 2013).

fundamental space that is nothing like the space of our ordinary three-dimensional experience: ray substantivalism posits a space of Points isomorphic to Hilbert space and wavefunction monism posits fundamental configuration space. Moreover, neither view posits fundamental local beables. The challenges facing monistic ray substantivalism and configuration-space-realist wavefunction monism are thus rather similar: it must be shown how the manifest image can be recovered from a fundamental picture that is quite unlike the manifest image.

The question of whether a fundamental ontology for non-relativistic quantum mechanics is capable of recovering the manifest image is part of a broader family of questions known as the 'quantum measurement problem': roughly, the challenge of providing a picture of fundamental physical reality according to quantum theory that is at least consistent with our total macroscopic evidence, including the behavior of measurement devices.[25] As I said earlier, monistic ray substantivalism entails all relevant propositions about the physical field that the wavefunction monist regards as fundamental. Therefore, if wavefunction monism succeeds in recovering our total macroscopic evidence in fundamental terms, then so does monistic ray substantivalism.

Third, I'd like to address a few objections that might be levelled against the kind of ontology articulated in this chapter. Although views along the lines of ray substantivalism have occasionally been mentioned in the literature, nothing close to a detailed statement of such an ontology has yet been given. One reason seems to be that views of this sort have been regarded as "unmotivated".[26] As I hope to have shown in this paper, this is untrue: ray substantivalism is the natural result of ontological theorizing that seeks to vindicate the paradigm shift towards symmetry techniques by giving symmetry fundamentalism a precise metaphysical underpinning. To be sure: ray substantivalism departs significantly from more familiar ways of thinking about the ontology of quantum mechanics. It is therefore no surprise that the prospect of a view along the lines of ray substantivalism has been dismissed as

[25](Wallace, 2008; Myrvold, 2018).

[26](Wallace, 2013, p. 216); (Wallace, 2017, p. 9). By contrast, (Ladyman, 2010, p. 157) and (French, 2013, p. 86) express favorable opinions about such views. Cf. (Bueno, 2011, p. 93n13).

"bizzare" in some parts of the literature.[27] But it should come as no surprise that metaphysical theorizing about fundamental physics might fail to validate the pre-theoretical intuitions that are informed by our exposure to the macroscopic realm—especially when it comes to a notion so far removed from our everyday experience as that of a fundamental state-space symmetry.

One more specific worry, due to Jill North, is that a fundamental space corresponding to abstract Hilbert space does not have enough structure from which "to recover the ordinary world of our experience."[28] The thought is that "Hilbert space does not support an objective, structural distinction between positions and other physical properties, like spin, in the way that the wavefunction's space do."[29] Translated to the present setting: the fundamental space of Points, together with its fundamental structure, privileges neither the normalized position-space nor the normalized momentum-space fields.

Although it is true that the space of Points does not support this kind of privileging, it does not follow that ray substantivalism does not have the resources to recover the manifest image—or at least, it does not follow that ray substantivalism fares worse on this front than (for example) configuration-space wavefunction monism. First of all, it would be incorrect to say that ray substantivalism has *too little* structure: as we have seen, ray substantivalism entails every proposition taken as fundamental by the wavefunction monist. If anything, one might think that ray substantivalism has a *surplus* of physical structure: in addition to the propositions about the position-space field taken as fundamental by the wavefunction monist, ray substantivalism also entails every proposition about momentum-space fields statable in the Hilbert space framework.

However, second of all, it would be a mistake to think that physical structure corresponding to momentum variables is surplus structure for either the wavefunction realist or the ray substantivalist. Although the configuration-space wavefunction realist describes the fundamental field in the configuration-space basis, she still requires conjugate position *and*

[27](Hawthorne, 2010, p. 145).
[28](North, 2013, p. 192).
[29]ibid.

momentum operators: for example, to derive the Heisenberg uncertainty relations, or to define the momentum wavefunction in terms of the Fourier transform of the position wavefunction. Similarly, as explained earlier, the ray-substantivalist definition of the configuration-space field of some State essentially involves the notion of momentum-almost-sameness: namely, via the definition of a simultaneous position-eigenpoint. Momentum structure is therefore not ontologically redundant in ray substantivalism, either.

Finally, let me briefly comment on the seemingly modal nature of vector and ray substantivalism. Recall: in physics, states are introduced as exhaustive and maximally specific characterizations of the fundamental properties of the material constituents of the world. As I briefly discussed in chapter 2, states therefore appear to represent instantaneous ways the world *could* be at an instant—much like possible worlds. By being committed to the claim that non-actualized states are every bit as real as the actualized one, vector and ray substantivalism might seem to entail a kind of modal realism.[30] But this impression is incorrect. At the risk of belaboring the point, according to vector or ray substantivalism, states aren't maximally specific and exclusive characterizations of anything and so are quite unlike possible worlds. Both views conceive of the space of Points as a fundamental space, the fundamental arena in which the history of the world unfolds. Just like in the case of classical state-space substantivalism, if Points (or relevant mereological fusions of them) are deserving of being characterized in modal terms—e.g. by referring to unactualized states as 'merely possible' states—this is due to the basic law according to which the world can be in no more than one state at every instant of time.[31] As in the classical case, neither Points nor the mereological fusions that play the role of states in ray substantivalism have intrinsic modal properties.

In this chapter, I developed ray substantivalism, an ontology for non-relativistic quantum mechanics that vindicates symmetry fundamentalism. According to this ontology, state-space symmetries correspond to fundamental relations, whereas the entities we would ordinarily

[30] I thank Dave Baker for pressing me on this point.

[31] Cf. the discussion at the end of chapter 2.

regard as the fundamental building blocks of the world are ontologically derivative of these relations.

Ray substantivalism is a novel ontology for non-relativistic quantum mechanics, one that is inspired by mathematical structures that have been and continue to be at the core of our best theories of fundamental physics. Although this is a significant result of its own, the explanatory power and theoretical utility of symmetry-fundamentalist ontological theorizing is likely to become fully transparent only once it is applied to the more sophisticated settings of quantum field theory and quantum gravity. For example: a symmetry-fundamentalist approach promises a metaphysical underpinning of the so-called 'classification of elementary particle kinds' in terms of the irreducible representations of Poincaré group as well as similar schemes involving the specification of color and flavor charge in terms of the group $SU(3)$. It also provides a blueprint for the development of a fundamental ontology that details the physical counterparts of spin networks in loop quantum gravity. Although the scope of this chapter was limited to non-relativistic quantum mechanics, its results are an important basis for future work on the symmetry-fundamentalist metaphysics of our best quantum theories.

Chapter 4

The Ontology of Symmetry Groups

The last two chapters illustrated two case studies that make good on the symmetry-fundamentalist claim that symmetries correspond to fundamental aspects of reality, whereas the material constituents of the world are non-fundamental and ontologically derivative on those aspects. The goal of this chapter is to illustrate what I take to be one of the most attractive features of symmetry-fundamentalist ontological theorizing: that it allows us to provide hitherto unavailable explanations of physical facts that are crucial to how our world works.

The starting point is an observation about the inferential structure of modern physics: important physical facts are inferred from claims about symmetry groups. I'll say more about this notion later; for the moment, think of a group as a mathematical tool that allows us to systematize and study the properties of entire classes of possible changes to the world—such as uniform spatial shifts or rotations. The group-based inferences I have in mind are united by their subject matter: they tell us *how many* values fundamental physical properties have. For example, the range of mass values can be captured in terms of the real line: there's a value of mass between $1kg$ and $5kg$ for every number in the interval $[1, 5]$ of the real line. By contrast, spin values differ by multiples of $1/2$, so for any finite interval of the real line, there's only a finite number of spin values. These features of mass and spin are derivable

from features of symmetry groups. For example, the discreteness of spin follows from a claim about the *rotation group*—the mathematical object that physicists use to systematize and study the general properties of rotations. For present purposes, all you need to know about this inference is that its premise attributes a certain mathematical property to the rotation group.[1] Group-based inferences of this kind are pervasive in particle physics and are responsible for some its most spectacular empirical successes, including the prediction of elementary particle kinds.[2]

Physicists have suggested that group-based inferences should be regarded as physical explanations.[3] Inferences of this sort do not merely provide convenient and useful mathematical characterizations of the range of possible values of physical properties; they tell us *why* physical properties have the values that they do, what *makes it the case* (for instance) that there are no more values of spin than there are non-negative natural numbers. This is a tempting thought: when claims about an entity are instrumental in deriving important physical facts, we are used to entertaining the hypothesis that this entity is a physical explainer of these facts. Such a hypothesis would be a particularly powerful way to vindicate its inferential role:

[1]Here are some further mathematical details. The fact that spin is discrete is a consequence of the fact that the rotation group (also referred to as $SU(2)$) is compact, together with what's known as the Peter-Weyl Theorem:

Theorem. If G is a compact group, then every irreducible representation of G is finite-dimensional, and every unitary representation of G is a direct sum of irreducible representations. Proof: Folland (2015, p 136).

The derivation of the fact that mass is continuous is part of the classification of elementary particle kinds in terms of the Poincaré group, a scheme inaugurated by Wigner (1939).

[2]Other physical properties whose values ranges are fixed by symmetry groups are isospin (derived from claims about $SU(2)$), flavor and color charge (both derived from claims about $SU(3)$), and helicity (derived from claims about the Poincaré group). For a survey of the rich history of symmetry techniques in 20th century, see (Martin, 2003).

[3]According to particle physics, each kind of elementary particle—electrons, quarks, and so on—corresponds to a specific combination of values of the fundamental physical properties mass, spin, electric charge and color charge. For example, electrons are negatively charged, colorless particles with spin 1/2 and a mass of $0.51099895 MeV$. The fact that the range of possible values of these properties can be derived from symmetry groups is regarded as a substantive claim about the explanatory hierarchy between particulate matter and symmetry groups. For example, according to Stephen Weinberg (1987, pp. 79-80), physical properties like mass and spin reflect "the way particles behave when you perform various symmetry transformations on them [...]" in a way that allows us to explain features of these properties purely in terms of symmetry groups: "Matter itself dissolves, and the universe itself is revealed as one large reducible representation of the symmetry group of nature." In a similar vein, Werner Heisenberg (1976, p. 924) claims that "the concept of a fundamental particle" should be replaced by "the concept of a fundamental symmetry". As Ne'eman and Sternberg (1991, p. 327) put it, "an elementary particle 'is' an irreducible representation of the group, G, of symmetries of nature".

symmetry groups are useful in inferring the relevant physical facts precisely because they occur in physical explanations of these facts, and so are involved in a full account of what makes it the case that these facts obtain. Moreover, group-based reasoning seems to explain facts that one might ordinarily have been tempted to regard as fundamental. Explaining why mass is continuous or why spin is discrete is a bit like explaining why space and time are continuous. The questions answered by putative group-based explanations belong to a family of questions concerning the nature of the physical world at the deepest level.

But there is something quite puzzling about the claim that group-based inferences are physical explanations. On the face of it, we are told that physical facts (such as the discreteness of spin) are explained by mathematical properties of mathematical objects: symmetry groups. But explaining why spin is discrete in terms of a purely mathematical fact about the rotation group seems no better than explaining facts about the structure of time in terms of facts about the real line. Nobody would explain the continuity of time by saying that time is measured by the real line, and the real line is continuous, so time is continuous. Clearly, there's a difference between the mathematical objects used to describe the world and the aspects of the world described in terms of them. If group-based reasoning is used to state physical explanations, we should not think that these explanations involve symmetry groups, but rather the physical entities to which symmetry groups correspond.

Unfortunately, physicists say hardly anything about what these physical entities are. As far as symmetry groups are concerned, there is a gap in the picture of the world described by physics. To fill this gap, we need to identify the physical entities to which symmetry groups correspond.

The goal of this chapter is to do just that. On the assumption *that* the relevant group-based inferences are to be read as physical explanations, I will explain *how* they can be read in this way—as explanations that appeal to facts about physical entities. I'm going to do this by articulating *group relationalism*, a hypothesis according to which the ontological counterparts of symmetry groups are four-place relations among the possible states of the world. This

view provides a precise sense of what it is to be realist about a symmetry group and therefore allows us to understand the group-based inferences given by physicists—those that proceed in terms of purely mathematical facts—as stating physical explanations.

There is an influential intellectual tradition in the philosophy of science that has been inspired in large part by the importance of symmetry groups in physics: ontic structural realism. On this view, the central ontological category in terms of which we should theorize about physical reality is that of 'structure', and symmetry groups are taken to characterize an important member of this category.[4] However, the ontic structuralist project is controversial. A common criticism is that structuralists have not sufficiently clarified the notion of structure.[5] It is therefore worth investigating whether physical explanations stated in terms of symmetry groups can be put on sound ontological footing without committing to ontic structural realism.

Here is my plan for this chapter. I introduce the target mathematical notions in section 4.1 and then develop group relationalism in section 4.2 and apply this view to group-based explanations in physics. Finally, section 4.3 investigates the vision according to which the ontological counterparts of symmetry groups are *fundamental*—one that has been suggested by physicists and philosophers alike—and sketches the relationship of this chapter to the symmetry-fundamentalist views developed in chapters 2 and 3.

4.1 Groups and Group Actions

There are two mathematical notions I need to introduce. The first notion is that of a group. Consider a weather vane that can rotate freely about its axis. What I'd like you to focus on is that rotations of the vane have the following four features. First, carrying out any two rotations successively is tantamount to carrying out some single rotation. For example: rotating clockwise by 45° followed by a clockwise rotation by 45° is tantamount to performing

[4](Ladyman, 1998; Castellani, 1998b; Brading and Castellani, 2003; Kantorovich, 2003; Ladyman et al., 2007; Roberts, 2011; French, 2014; McKenzie, 2014a; Brading et al., 2017).

[5]For example, (Dorr, 2010; Wolff, 2011; Arenhart and Bueno, 2015); cf. (McKenzie, 2017).

a single clockwise rotation by 90°. Second, every rotation can be reversed by what's called the *inverse* of that rotation. For example: a clockwise rotation by 90° can be reversed by a clockwise rotation by 270°. Third, there is a rotation that amounts to doing nothing, the *identity element* among the rotations: the rotation by 0°. Moreover, the order in which rotations are combined doesn't matter, in the following sense: the outcome of performing the combination of clockwise rotations by 45° and 90° followed by a clockwise rotation by 180° is the same as first rotating clockwise by 90° and then performing the combination of clockwise rotations by 45° and 180°.

These features of rotations are independent of the object we're considering. Rotations of a football or a brick about a fixed axis have the same properties as rotations of a weather vane. In fact, these four properties are common to rotations of *any* object about a fixed axis. Intuitively, we can think of the rotation *group* as a mathematical object that captures what is common to rotations of any object whatsoever.

There are two components to the group of rotations of the vane. First of all, a set whose members we each think of as a possible rotation of the vane. Elements of the rotation group are in one-to-one correspondence with angles—one for each real number in the open interval $[0, 360)$. We also lay down a binary function on that set, the function that captures what I above referred to as 'combining' two rotations. We can model this function as the addition of angles (modulo 360). The set, together with a binary function on it, is called a *group* just in case it has the four properties detailed in the previous paragraph, also known as the *group axioms*.[6] Abstractly, a group is just a set with a binary operation that satisfies these axioms.

[6]In technical terms: a group is a set G to gether with a map $\cdot$ on G satisfying the following conditions:

1. Closure: $g \cdot g' \in G$ for all g, g';
2. Inverse element: for every $g \in G$ there is an element (called the *inverse element*) $g^{-1} \in G$ such that $g \cdot g^{-1} = g^{-1} \cdot g = e$';
3. Identity element: there is an element (called the *identity element*) $e \in G$ such that $e \cdot g = g \cdot e = g$ for all $g \in G$;
4. Associativity: $g \cdot (g' \cdot g'') = (g \cdot g') \cdot g''$, for all $g, g', g'' \in G$.

The second notion I need to introduce is that of a group *action*. Physics is not primarily concerned with the notion of a group in abstraction, but rather with the way in which groups can be used to capture possible changes to the physical world. Indeed, the group-based inferences we're concerned with here don't involve groups by themselves, but rather groups in conjunction with a specification of a possible change of the physical world (such as a rotation) for each group element. This is specification is called a group action. I said above that the rotation group can be thought of as the mathematical notion that captures what rotations of different objects have in common. By contrast, think of an action of the rotation group as what rotations of different objects disagree about.

Let's make this a bit more precise. The kinds of objects on which group actions in physics are defined aren't physical objects like weather vanes, but rather *mathematical spaces* that encode salient aspects of those objects. For example, consider the space whose points correspond to possible orientations of the weather vane—i.e. the space that consists of one point at which the vane points north, one where it points east, and so on. An action of the rotation group on this space assigns a *permutation* of the space to each group element: a bijective function that takes any given point to its image under a rotation by some fixed angle. For example, the permutation assigned to a clockwise rotation by 90° takes the point at which the vane is directed north to the point at which it is directed east, and the point at which the vane is directed east to the point at which it is directed south, and so on. For this assignment of permutations to group elements to be an *action* of the rotation group, the composition of permutations must 'respect' the composition of abstract group elements, in the sense that the permutation assigned to the sum $\alpha + \beta$ of any two angles α and β must be identical to the function composition of the permutation assigned to α and the permutation assigned to β. Moreover, the identity element of the group (represented by 0°) must be assigned the identity permutation: the function that maps every point to itself.[7]

[7]In technical terms, an action of a group G on a space X is a map $\phi : G \times X \to X$, $(g, x) \mapsto \phi(g, x)$ such that

1. $\phi(g, \phi(g', x)) = \phi(g \cdot g', x)$

There are two different kinds of group actions that are commonly found in physics, and we'll see shortly that the difference between them is important for the purposes of ontological theorizing about groups. So I should explain what these kinds of group actions are.

Group actions can be rather trivial. For example, an assignment of the identity permutation of the space of orientations of the weather vane to *every* group element (i.e. to every angle) counts as an action of the rotation group. But insofar as we'd like to use the notion of a group to capture possible changes to a physical system, we're interested in actions of a less trivial sort: we want actions that assign *distinct* permutations to *distinct* group elements. Group actions of this sort are called *faithful*.[8] The action of the rotation group we considered earlier is of this kind. The permutation which assigns to every point its image under a rotation by 90° is distinct from the permutation that assigns to every point its image under a rotation by 45°, and so on.

The second kind of group action is a subclass of the faithful actions. Observe that some faithful actions have *fixed points*: points not moved around by some permutation distinct from the identity permutation. To illustrate this, let's abstract slightly from the example of a weather vane and consider instead rotations about a point in the Euclidean plane. (Think of this as the point where the weather vane is attached to its axis.) This action associates with 90° the permutation that maps every point on the positive y-axis to a unique point on the positive x-axis and every point on the negative y-axis to a unique point on the negative x-axis; with 180° the permutation that maps every point on the positive y-axis to a unique point on the negative y-axis and every point on the negative y-axis to a unique point on the positive y-axis, and so on. As before, this action assigns distinct permutations to every group element, so it's faithful. However, each permutation does not move around the point

2. $\phi(e, x) = x$

for all $g, g' \in G$ and $x \in X$, where e is the identity element in G. The first condition ensures that ϕ respects the multiplication operation on G; the second ensures that the identity element gets mapped to the identity permutation of X.

[8] In technical terms: an action $\phi : G \times X \to X$ is faithful iff for all $g \neq g'$ in G there is an a in X such that $\phi(g, a) \neq \phi(g', a)$.

about which we're rotating. This point is fixed by every group element; it is a *fixed point* for every group element. Group actions for which no group element has a fixed point are called *free* actions.[9] An important range of group actions in physics are of this kind: those that correspond to shifts in physical space or space-time. For example, shifts in the Euclidean plane can be characterized in terms of a free action of $\mathbb{R}^2$ (a group with regard to vector addition): the action that assigns to each vector $v \in \mathbb{R}^2$ the permutation that takes x to $x+v$, for any point x in the Euclidean plane. Similarly, space-time shifts in relativistic physics can be characterized in terms of a free action of $\mathbb{R}^4$.

These are all the technical notions we need to provide an account of the candidate ontological counterparts of group actions. That's what I'll do now.

4.2 Group Relationalism

Group actions are mathematical notions that are defined within the mathematical apparatus used to state physical theories. If we are to understand group-based inferences as stating physical explanations, we need an account of the entities that are candidates for being the ontological counterparts of groups. The goal of this section is to develop such an account.

Let's start with a tempting but ultimately inadequate idea. Suppose we're already committed to the existence of a space X of physical points, and suppose that we can characterize X in terms of a mathematical space on which a group G acts faithfully; call this action ϕ. Then there is a fairly straightforward sense in which we can understand this action as specifying a collection of *dyadic relations* on X. Every permutation that ϕ assigns to a group element corresponds to a set of ordered pairs of points in X: the set consisting of ordered pair $\langle a, b\rangle$ just in case that permutation maps a to b. Since ϕ is faithful, there is a distinct set of ordered pairs of points in X for each group element. The idea is then to posit a dyadic relation on X—one for each group element—that relates points a and b just in case

[9] In technical terms: an action $\phi : G \times X \to X$ is free iff for all g in G and some a in X, $\phi(g, a) = a$ implies $g = I$.

the corresponding permutation maps a to b. Moreover, since the permutations assigned to the group G by ϕ have the structure of a group with regard to function composition, the same is true for the corresponding dyadic relations with regard to the appropriate relational analogue of function composition.[10]

But this proposal is somewhat dissatisfying. All we seem to accomplish is to replace permutations with dyadic relations. This feels less informative than we might have hoped: we already knew that both permutations and dyadic relations are associated with sets of ordered pairs. Moreover, on this proposal, we posit as many dyadic relations as there are group elements: for continuous groups, a continuum infinity!

But we can do better. As I will show now there is a way to capture the structure of a group action on the mathematical surrogate of X just in terms of a single four-place relation among points of X.

To illustrate this idea, let's go back to the example of shifts of points in the two-dimensional Euclidean plane; that is, an action of the shift group $\mathbb{R}^2$ on this space. For any vector in $\mathbb{R}^2$, this action tells us not only how to shift some specific point by that vector, but *all* points in the Euclidean plane. This brings us to an idea that is central to the view I'll develop: the action of the shift group gives us a notion of shifting *different* points by the *same* amount. That is: given that point b is the outcome of a shift of a, we can ask: what is the outcome of doing to another point c what we did to a in order to obtain b? More precisely: the notion of

[10]Here is how the 'composition' of dyadic relations can be defined. Let $\mathcal{G}$ be the set of dyadic relations on X such that there is a one-to-one correspondence between elements of the group G and relations in $\mathcal{G}$. Then, for any $A, B, C \in \mathcal{G}$, C is the composition of A and B—or $C = A \cdot B$, for short—iff the following conditions are met:

1. $\forall x \forall y \exists w \exists z \exists w' \exists z' (Cxy \wedge Awz \wedge Bw'z' \rightarrow z = y \wedge w' = x \wedge z' = w)$;
2. $\forall w' \forall z' \exists w \exists z \exists x \exists y (Bw'z' \wedge Awz \wedge Cxy \rightarrow z = y \wedge w' = x \wedge z' = w)$;
3. $\forall w \forall z \exists w' \exists z' \exists x \exists z (Awz \wedge Bw'z' \wedge Cxy \rightarrow z = y \wedge w' = x \wedge z' = w$.

The first condition says that for any points x, y related by C there is a point z between x, y such that B relates x and z and A relates z and y. The second condition says that, for any points w', z' related by B, there is a point z to which A relates z' and to which C relates w'. The third condition says that, for any points w, z related by A, there is a point w' that is B-related to w and that is C-related to z.

The group axioms can then be specified in a way that's structurally similar to their mathematical counterparts; see footnote (6).

a group action allows us to say that points a and b differ *as much* by a shift as do points c and d. This is a four-place relation among points of Euclidean space. Intuitively, this relation captures whether the differences between any two points are the same as the differences between any two other points; I'll call it the *same-difference relation*. More precisely, an action ϕ of a group G on the mathematical surrogate of a space X of physical points specifies the following relation among points of X:

> **Same-Difference Relation.** A four-place relation. Relates four points a, b, c, d in X iff there is a group element $g \in G$ such that the permutation assigned to g by ϕ maps a to b and c to d. Write $S_{G,\phi}abcd$.

Group relationalism is the view that the ontological counterpart of a group action on the mathematical surrogate of a physical space X is the same-difference relation corresponding to this action. To show that the notion of a same-difference relation has what it takes to satisfy this role, we need to demonstrate that the structure of the relevant types of group actions can be captured entirely in terms of same-difference relations. If successful, this will show that we don't need to posit an infinity of dyadic relations to make sense of group-based inferences as physical explanations: a single four-place relation is enough.

My goal in the remainder of this section is to explain how group structure can be encoded in same-difference relations. More specifically, the challenge is as follows. Let ϕ be an action of G on the mathematical surrogate of X. Recall: each permutation that ϕ assigns to an element of G corresponds to a set of ordered pairs of points in X: the set consisting of $\langle a, b\rangle$ iff this permutation maps a to b. Now, let $\mathbb{G}$ be the set consisting of all such sets of ordered pairs: one for each permutation assigned to elements of G by ϕ. (For faithful actions, members of $\mathbb{G}$ are in one-to-one correspondence with group elements.) To show that the same-difference relation corresponding to ϕ captures the structure of G, we need to demonstrate how $\mathbb{G}$ can be defined just in terms of the same-difference relation specified by ϕ. This is what I'll do now.

4.2.1 Free Actions

Free actions allow for a simple and uniform definition of the sets of ordered pairs in terms of the relevant same-difference relation. The key observation is that, if ϕ is a free action on the mathematical surrogate of X, the corresponding same-difference relation on X is an equivalence relation among ordered pairs of points in X, and $\mathbb{G}$ is just the set of equivalence classes of this relation.

Let me explain. If ϕ is free, then no two points in the mathematical surrogate of X are related by more than one group element. In other words: if a is mapped to b by some group element, then this is the only group element that maps a to b.[11] For example, for any two points a, b related by a shift in the Euclidean plane, there is *exactly one* shift that maps a to b.[12] This is sufficient for showing that the corresponding same-difference relation is an equivalence relation. It is reflexive and symmetric by construction: clearly, if a, b differ by some shift, then a, b differ by the same shift as a, b; and if a, b differ by the same shift as c, d, then c, d differ by the same shift as a, b. Slightly less obviously, the same-difference relation is also transitive. Suppose that a, b are related by the same shift as c, d, and c, d are related by the same shift as e, f. Since the action of the shift group on Euclidean space is free, there is a unique shift that relates c, d. This entails that a, b are related by the same shift as e, f. Moreover, this shows that the equivalence classes of ordered pairs of points in X under the relevant same-difference relations each correspond to a single group element. Thus, if ϕ is a free action of G on the mathematical surrogate of X, the set $\mathbb{G}$ is just the set of equivalence classes of $S_{G,\phi}$ on X.

As I will now argue, this strategy cannot be applied to group actions that are faithful

[11] More precisely: if ϕ is free and for any two points a, b in the mathematical surrogate of X, $b = \phi(g, a)$ and $b = \phi(g', a)$, then $g = g'$. The proof is as follows. Let ϕ be free and suppose that there are group elements $g \neq g'$ such that for some point a, $\phi(g, a) = \phi(g', a)$. If $g' = I$, it is immediate that a is a fixed point for g. Otherwise, apply a permutation by g'^{-1} to obtain $\phi(g'^{-1}, \phi(g, a)) = \phi(g'^{-1}, \phi(g', a))$. By the definition of a group action, $\phi(g'^{-1}g, a) = \phi(I, a) = a$. So, a is a fixed point for $g'^{-1}g \neq I$.

[12] Of course, this does not mean that there isn't some combination of other shifts that will also to take a to b. By the group axioms, this combination of shifts—the permutation associated with this combination of shifts—will be identical with the permutation that takes us from a to be b 'directly', so to speak.

but not free. This is a problem. Non-free faithful actions are the basis of an important range of the group-based inferences we are concerned with—such as the derivation of spin's discreteness in terms of the rotation group. If group relationalism is to keep its promise of allowing us to understand such group-based inferences as physical explanations, it should be applicable to actions of this kind. I'll now explain how the group relationalist can accomplish this.

4.2.2 Non-Free Faithful Actions

If an action is faithful but not free, the corresponding same-difference relation is not an equivalence relation.[13] Recall our earlier example of rotations about a point x in the Euclidean plane. This is a faithful action of the rotation group in two-dimensions, a group known as SO(2). Call this action ϕ_x. The same-difference relation $S_{SO(2),\phi_x}$ corresponding to this action holds among points a, b, c, d in the Euclidean plane iff there is some rotation about point x that maps a to b and c to d. Now, consider two distinct group elements—say, those corresponding to clockwise rotations by 45° and 90°—and suppose that the permutation associated with 45° maps a to b and the permutation associated with 90° maps c to d, where a, b, c, d are all distinct from x. Since x is a fixed point under every rotation, it follows that both permutations map x to itself. So, $\langle a, b\rangle$ and $\langle c, d\rangle$ each bear the same-difference relation to $\langle x, x\rangle$ but not to each other. This illustrates that same-difference relations corresponding to non-free faithful actions are not transitive and so aren't equivalence relations.

Faithful actions require a more sophisticated approach. To begin with, note that the reason why ϕ_x isn't free is the existence of a fixed point, x. By the same token, the failure of transitivity of $S_{SO(2),\phi_x}$ is due to the fact that *every* ordered pair of points related by some rotation about x bears the same-difference relation to $\langle x, x\rangle$.

But note that ϕ_x is 'almost' free: for any two points a, b distinct from x, if a and b are related by the permutation assigned to (say) a clockwise rotation by 90°, then the

[13] I am indebted to Neil Dewar for helpful discussions on this point.

permutation assigned to this angle is the *unique* permutation that maps a to b.[14] This suggests the following modification of our strategy. Recall: the goal is to define the set of sets of ordered pairs $\mathbb{SO}(2)$ just in terms of the same-difference relation $S_{SO(2),\phi_x}abcd$. We start out by defining a four-place relation S_x that obtains between points a, b, c, d iff $S_{SO(2),\phi_x}abcd$ and a, b, c, d are all distinct from x. The key feature of this relation is that, unlike $S_{SO(2),\phi_x}$, it never applies to x. This means that S_x is an equivalence relation on sets of ordered pairs. Next, let the *S-extension* of an equivalence class e of S_x be the set e' such that e' contains an ordered pair iff it bears $S_{SO(2),\phi_x}$ to *every* ordered pair in e. An equivalence class of S_x differs from its S-extension in that the latter, but not the former, contains $\langle x, x\rangle$. $\mathbb{SO}(2)$ can therefore be defined as the set of S-extensions of the equivalence classes of S_x.

We can take a similar approach to other faithful non-free actions. By way of illustration, I want to focus on the case of the group SO(3) of rotations about a point in three-dimensional Euclidean space. This is an important case study because SO(3) plays an important role in the derivation of the fact that spin is discrete, as well as in the classification of elementary particle kinds according to the Poincaré group.[15]

Here is how we can apply our strategy to SO(3). Consider an action ρ_x of this group as rotations about a point x in three-dimensional Euclidean space. You can think of this action as consisting of rotations of Euclidean space about arbitrary axes intersecting x: for each element R in SO(3), there is some axis through x and some angle such that the permutation associated with R corresponds to a rotation by that angle about that axis. This action is faithful, but not free: just like in the case of SO(2), x is a fixed point of every permutation. Unlike the action of SO(2) considered earlier, however, x is not the only fixed point. For every group element R, the points on the axis of rotation of R are fixed by R. As a result, the same-difference relation $S_{SO(3),\rho_x}$ is not an equivalence relation even when restricted to

[14] Of course, a counterclockwise rotation by 270° will also map a to b; but this is because the action of SO(2) assigns the same permutation to clockwise rotations by 90° that it assigns to counterclockwise rotations by 270°.

[15] The group that is involved in the derivation of spin's discreteness is called the 'double cover' of SO(3), also known as $SU(2)$. In the classification of elementary particle kinds, the double cover of SO(3) occurs as the 'little group' of the Poincaré group for particles with positive mass; see (Kim and Noz, 2012).

points distinct from x.

This indicates that we need to modify the restriction we imposed in the two-dimensional case. We need to remove not just $\langle x, x\rangle$ from the scope of the same-difference relation, but *all* ordered pairs consisting of the same point: ordered pairs of the form $\langle a, a\rangle$ for any a. We can achieve this by defining a relation that relates an ordered pair only if each of the points in this pair are distinct. More precisely, let $S_{\neq}$ be the relation that holds between points a, b, c, d iff $S_{SO(3),\rho_x}abcd$ and $a \neq b$ and $c \neq d$. The latter two conjuncts ensure that $S_{\neq}$ is an equivalence relation: whenever $S_{\neq}abcd$, there is exactly one R in SO(3) such that the permutation assigned to R takes a to b and c to d. As before, we obtain a range of S-extensions of the equivalence classes of $S_{\neq}$, each of which corresponds to some element in SO(3). But unlike the two-dimensional case, there is no equivalence class of $S_{\neq}$ whose S-extension corresponds to the identity element of SO(3), for the simple reason that $S_{\neq}$ by definition does not relate ordered pairs consisting of the same point. What's missing, in other words, is the *identity set* of Euclidean three-space: the set that consists of ordered pairs like $\langle a, a\rangle, \langle b, b\rangle$, and so on. The following definition reflects this observation: $\mathbb{SO}(3)$ is the union of the set of S-extensions of the equivalence classes of $S_{\neq}$ with the identity set of Euclidean three-space.

This modified strategy also delivers the right result for rotations in the Euclidean plane. The main difference is that, since the restriction imposed by the constrained relation applies to all ordered pairs that involve the same point, it excludes more ordered pairs than it needs to. Nonetheless, $\mathbb{SO}(2)$ can be defined as the union of the relevant S-extensions with the identity set.

The result of this section is therefore that the group relationalist has the resources to capture the structure of a faithful action of the rotation group just in terms of the corresponding same-difference relation. We have seen, moreover, that there are significant differences between the group-relationalist account of free and faithful group actions. Faithful actions require the definition of a restriction on same-difference relations, whereas free actions do not.

Moreover, we saw that those restrictions must be responsive to the specific way in which a faithful action fails to be free. Whereas the group-relationalist approach to free actions works for any free action whatsoever, there seems to be no uniform group-relationalist strategy for faithful actions.

Nonetheless, the discussion in this section suggests a general recipe for how group structure can be obtained from the same-difference relation in specific cases. The key is to find the appropriate constraint for the same-difference relation corresponding to this action, a constraint whose job it is to exclude the points fixed by some group element. It is important that this constraint excludes the points in a way that does not require a one-by-one specification of each group element: such an account would potentially require an infinity of constraints—one for each group element—and so would not fare much better than the proposal, briefly discussed earlier, that posits one dyadic relation for each group element. Once this constraint is identified, the set of ordered pairs corresponding to the relevant group can then be defined in terms of the resulting set of S-extensions.

So far, I've introduced and explained group relationalism, a hypothesis about the physical entities that are candidates for being the ontological counterparts of symmetry groups. But we're not done yet. We want to deploy same-difference relations to understand inferences on the basis of symmetry groups as physical explanations.

As I argued in the introductory chapter, the relevant notion of symmetry is that of a *state-space symmetry*: a state-space automorphism that preserves the characteristic structure of state space. Indeed, the only symmetry-related premise in the derivation of the fact that spin is discrete is that the rotation group acts in terms of state-space symmetries—i.e. in terms of unitary transformations. The same is true for other relevant symmetry-based inferences, such as derivation of the structure of mass from claims about the Poincaré group.[16]

Call a group that acts in terms of symmetries on the mathematical surrogate of state space a *state-space symmetry group*. In quantum mechanics, state-space symmetries are called

[16]See footnote 1.

unitary transformations; in classical mechanics they are called symplectomorphisms.[17]

It is now a simple matter to deploy group relationalism in order to understand the mathematical derivations in terms of symmetry groups as physical explanations. If a derivation of some physical fact proceeds by appeal to a group G that acts in terms of state-space symmetries, the corresponding physical explainer is the same-difference relation on state space that is the counterpart of this action of G. For example: according to the group relationalist, the physical explainer of the fact that spin is discrete is the same-difference relation that corresponds to the relevant action of the rotation group in terms of state-space symmetries.

Let me briefly recapitulate what I have done so far. Our goal was to articulate an ontology that details the physical counterparts of the symmetry groups that figure in the relevant symmetry-based explanations. According to the group relationalist approach, state-space symmetry groups correspond to same-difference relations on state space: four-place relations that we can intuitively think of as capturing the idea of 'rotating by the same amount in the same direction'. Group relationalism therefore gives us a precise account of the physical entities that are appealed to in the symmetry-based explanations that establish (for example) the discreteness of spin or the continuity of mass: according to group relationalism, the relevant mathematical derivations state physical explanations in terms of the corresponding same-difference relations.

4.3 Discussion

The metaphysical vision that inspired this dissertation is one according to which symmetries correspond to fundamental aspects of reality, whereas the material constituents of the world are ontologically derivative on those aspects. So the reader may wonder how the results of this chapter hang together with the symmetry-fundamentalist views developed in the previous chapters. This is what I turn turn to now.

[17] Another name for state-space symmetry is *Wigner symmetry*, so-called in recognition of the mathematical physicist Eugene Wigner who spearheaded the study of unitary actions of symmetry groups on the state spaces of quantum systems; (Wigner, 1931, p. 251-3); cf. (Weinberg, 1995, p. 51).

First of all, this chapter showed that symmetry structure is not the same as symmetry *group* structure. It is one thing to posit fundamental relations that capture individualistic (particle) claims in terms of spatial and momentum shifts; it's another to capture the group structure of spatial and momentum shifts. The former task can be done in terms of two-place relations among states; the latter requires four-place relations of the sort introduced in this chapter.

Second, the claim that symmetries are fundamental is often expressed in conjunction with or in the guise of the claim that the *groups* which systematize these symmetries are fundamental. At this point, it should go without saying that the claim we're entertaining is not that the ontological inventory of fundamental physical reality involves state-space symmetry groups themselves, but rather that it contains the physical counterparts of these groups. Group relationalism gives precise content to this claim: fundamentalism about symmetry groups is the thesis that the relevant same-difference relations are fundamental.

Third of all, for reasons familiar from earlier parts of this dissertation, fundamentalism about state-space symmetry groups entails fundamentalism about state space—at least on a fairly straightforward view of what it is for an entity (such as an object, property, or relation) to be fundamental. Recall that, according to this view, for an entity to be fundamental is for it to figure in some fundamental propositions.[18] If a same-difference relation is fundamental, then this is presumably so because there is a fundamental proposition of the form $Sabcd$, where a, b, c, d are instantaneous states. This means that the states that are the relata of S are fundamental, too. On this view of entity fundamentality, the fundamentality of the same-difference relation on state space entails the fundamentality of state space. The claim that symmetry groups are fundamental therefore amounts to the claim that actions of state-space symmetry groups correspond to fundamental relations among fundamental entities: namely, same-difference relations among instantaneous states of the world.

Finally, there is a question about the relationship between the almost-sameness and the

[18] (Schaffer, 2009; Audi, 2012; Dasgupta, 2014).

same-difference relations. If we want to make sense of the claim that symmetry groups are fundamental while also making good on the promise that every proposition about the material constituents of the world can be reduced to a proposition about symmetries, do we need to posit both kinds of relations?

The answer is no: there is good reason to think that the almost-sameness relations can be incorporated into the same-difference relations. In the classical setting, for example, the idea is to posit four place relation of *almost same-difference*, a relation which obtains between four states a, b, c, d iff there is some particle such that a differs from b in a spatial/momentum shift of that particle as much as do c and d, and both a, b and c, d differ *only* in the position/momentum of that particle. The almost-sameness relations can then be defined in terms of this four-place relation: for example, states a, b are position-almost-same iff the almost same-difference relation holds of the quadruple a, b, a, b. This view deserves a more detailed investigation in its own right, one that is beyond the scope of this dissertation.

The concern of this chapter has been a powerful inferential strategy according to which facts about the possible values of physical properties (such as the fact that spin is discrete) are derived from claims about mathematical objects: symmetry groups. We started out with the observation that there is a gap in the picture of the world according to physics: although we would like to understand group-based inferences as physical explanations in terms of the purely physical entities that correspond to symmetry groups, physicists say hardly anything about what those entities are.

My goal in this chapter has been to fill this gap and thereby to allow us to regard the group-based inferences as physical explanations. More specifically, the goal has been to develop *group relationalism*, an ontological hypothesis according to which state-space symmetry groups correspond to same-difference relations on state space: four-place relations among states that capture the notion of (for example) 'rotating by the same amount in the same direction'. The main result of this chapter is therefore that derivations of the possible values of physical properties that proceed by appeal to state-space symmetry groups can be

understood as physical explanations in terms of the same-difference relations on state space that correspond to these groups.

Chapter 5

The Metaphysics of Invariance

In the previous three chapters, we took for granted a fairly general conception of a symmetry of an object (such as a mathematical space) as a transformation that leaves invariant certain salient features of that object. The goal of this chapter is to focus on the notion of invariance itself, and more specifically, on the notion of invariance that is involved in claims about the invariance of a *physical property* under symmetries.

We are familiar with the idea that a property is *modally* invariant under a symmetry: the appearance of a round dinner plate *would* be the same if it *were* rotated. But the notion of property invariance operative in fundamental physics is distinct from this modal notion: it is a notion that is thought to underwrite a range of claims concerning the ontological priority of symmetries over the material constituents of the world, such as particles. More specifically, property invariance is supposed to capture a correspondence between the essential properties of elementary particles and continuous symmetries of particle systems. The basic idea is that there is a notion of property invariance such that, for each continuous symmetry, there is a specific, limited collection of properties that is invariant under it. As Steven Weinberg puts it, each of the essential properties of elementary particles, such as mass and spin are each "associated with" specific continuous symmetries.[1] This link has played a key role in

[1]Weinberg (1992, 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

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

This suggests that the relevant notion of invariance—whatever it is—must be more fine-grained than modal invariance. The essential properties of particles are modally invariant under all the relevant symmetries, but they do not all count as invariant under these symmetries in the target sense. Indeed, modal invariance is quite permissive: for any continuous symmetry, there is a wide range of properties modally invariant under it. By contrast, the correspondence between a continuous symmetry and a physical property that the target notion of property invariance is supposed to capture is quite restrictive: only a sparse range of properties that are modally invariant under a symmetry also qualify as invariants in the target sense. For example, 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 modally invariant under rigid rotations about any given point.[3] And yet, for every point x, only the magnitude of total angular momentum about x is invariant under rigid rotations about x in this more fine-grained sense. Call this notion the *fine-grained invariance* of a property. The goal of this chapter is to give a precise and transparent account of this notion.

According to an influential line of scholarship, the link between properties and symmetries

symmetries was first hypothesized by Eugene Wigner (1939).

[2](Weinberg, 1992, p. 138n, emphasis added). Elsewhere, Weinberg claims that the relevant link 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, 1992, p. 139).

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

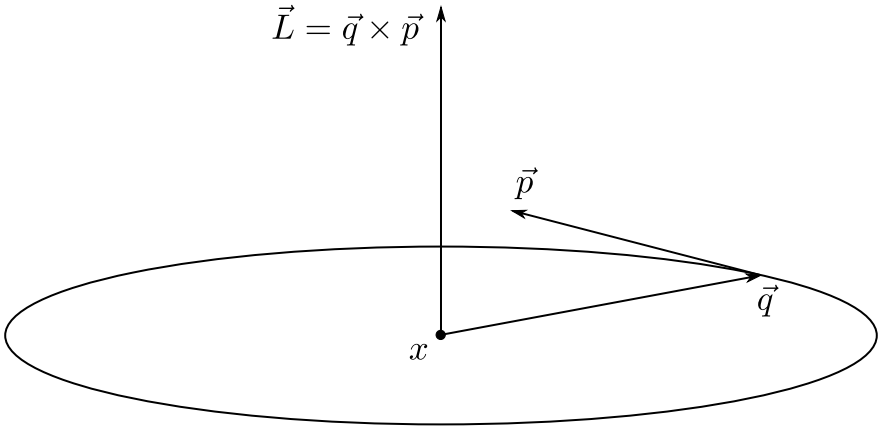


Figure 5.1: Illustration of a classical one-particle system in orbital motion about x. $\vec{q}$ represents the position of the particle in a co-ordinate system originating at x, $\vec{p}$ its linear momentum and $\vec{L} = \vec{q} \times \vec{p}$ its angular momentum about x.

captured by fine-grained invariance 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 link between physical properties and symmetries is "the basis of the claim that such properties should be conceived of structurally,"[4] 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]."[5] However, if we are to use claims about a link 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.

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

[5](McKenzie, 2014b, p. 21). This is echoed by other authors; e.g. Ladyman (1998); Ladyman et al. (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).

The correspondence between properties and symmetries captured by a putative notion of fine-grained invariance is distinct from a more familiar correspondence, captured by Noether's theorem, between continuous symmetries and *conserved quantities*.[6,7] 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 not the connection physicists have in mind when they claim that a given property is linked to a continuous symmetry in the relevant sense. Whereas this link concerns "the way particles behave when you perform various *symmetry transformations*" such as rotations,[8] Noether's theorem concerns the way particles behave *under temporal evolution*, i.e. whether certain physical quantities are conserved. The correspondence captured by Noether's theorem is therefore distinct from the link between properties and continuous symmetries established by the relevant notion of fine-grained invariance.

The question we're asking has a standard mathematical answer—one that physicists like Weinberg presumably have in mind when making the more colorful claims cited above. This answer goes roughly as follows. The notion of a continuous symmetry is mathematically formalized 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 every property that is a fine-grained invariant under a particular symmetry group is

[6](Noether, 1918).

[7]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.

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

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

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 translations.[11] The standard mathematical answer is thus that a physical property P is a fine-grained invariant under a continuous symmetry group G in the relevant sense 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 characterizing 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 characterize the notion of fine-grained invariance of a property under a symmetry group as the feature that property has just in case 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

[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 (possible worlds, facts, propositions, objects, properties, quantities).

[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.)

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? In other words: the goal is not to explicate the *mathematical* notion of a Casimir invariant, but rather to give an account of the *physical* notion to which it corresponds.

I proceed as follows. I explain the standard mathematical answer in section 5.1 by considering the case study of total angular momentum about a point x of a classical N-particle system and the group of rigid rotations of the system about x. Subsequently, section 5.2 introduces and explains the notion of a fine-grained invariant, the physical notion that corresponds to the notion of a Casimir invariant. This notion is 'fine-grained' in exactly the right way: not every invariant (in the ordinary modal 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. Finally, section 5.3 draws some contrasts with ontic structural realism.

5.1 Casimir Invariance

As noted in the introduction, the notion of a fine-grained invariant is mathematically characterized in terms of the notion of a Casimir invariant. The goal of this section is to introduce this notion and to characterize the theoretical and metaphysical framework I'll be using throughout this paper.

5.1.1 The Framework

To begin with, I need to tell you what notion of symmetry is operative in the claim that some property is a fine-grained invariant under some symmetry. As in the previous chapters, the relevant notion is that of a state-space symmetry.

The notion of a fine-grained invariant is of particular significance in quantum field theory—a framework whose mathematical complexity is substantial. Luckily, the fact that we're dealing with state-space symmetries means that we don't need let ourselves be sidetracked by these complexities: we can understand the notion of a fine-grained invariant within the relatively simple setting of N-particle Hamiltonian mechanics. This is the framework I'll be using in this chapter.[13]

To understand the notion of a Casimir invariant in the context of Hamiltonian mechanics, I will need to introduce this framework in a bit more detail. The phase space of an N-particle system in three-dimensional Euclidean space 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.[14] 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.

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 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.[15] For the sake of brevity, I will

[13]To be sure: N-particle Hamiltonian mechanics has no hope of being the true theory of fundamental physical reality. But that's fine: this paper is not intended as a contribution to the metaphysics of classical mechanics. Instead, the point of this investigation is to draw metaphysical conclusions that can be transposed to more cutting-edge physical theories, or that (at the very least) allow us to identify general strategies for handling such cases.

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

[15]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

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'.[16]

I will take for granted that physical quantities can be distinct despite being represented by the same phase space function. More precisely: call two quantities Q and 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: $f_Q = f_{Q'}$. Then, the claim I take for granted is 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 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*. This motivates

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 chapter, I follow this practice for the sake of brevity.

[16]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. Distinctions based on naturalness are familiar from the metaphysics of properties; cf. (Lewis, 1986, pp. 59).

the assumption that there are distinct state-space co-extensive quantities.

In this chapter, I focus 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* axis through x.[17]

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 exactly once 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 given rotational flow is

[17] 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.

interpreted as the corresponding angle of rotation; and states related by a rotation are generally not related by dynamical evolution, as illustrated in figure (5.2).

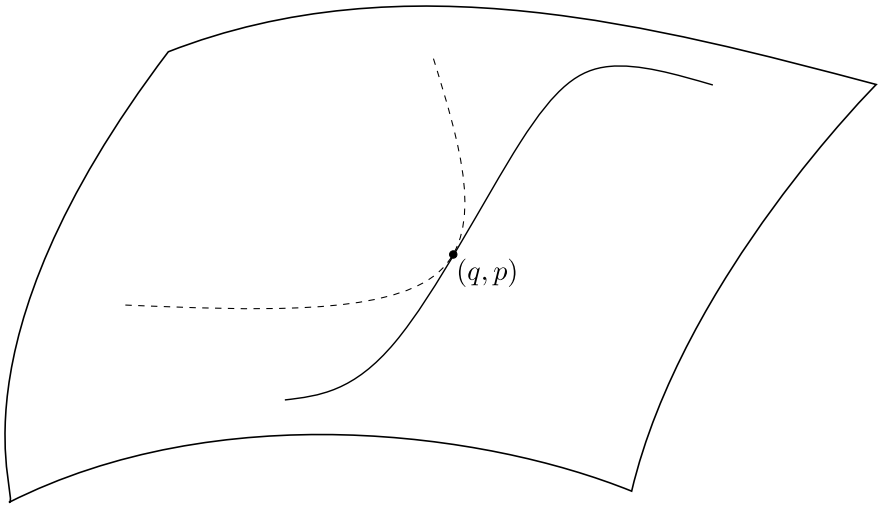


Figure 5.2: Illustration of phase space with a dynamical flow line and a rotational flow line (dotted) through phase space point (q, p).

This is where a crucial aspect of Hamiltonian mechanics comes into play: the phase space flows representing (for example) dynamical evolution and rotations about some axis in physical space are what's referred to as *Hamiltonian* flows. I'll explain this notion in a moment.

But first, let me tell you why it's important, in the present context, that these flows are Hamiltonian. The first reason is that Hamiltonian flows are *symmetries* of phase space: they preserve the symplectic form, the characteristic structural feature of state space. (Again, you can think of this as the claim that Hamiltonian flows preserve the volumes of phase space regions.[18]) Second, each of these flows has an important mathematical relationship to some specific smooth phase space function: each Hamiltonian flow is *determined* by such a function in a way that *preserves* this function. The function which determines a Hamiltonian flow is called the *generator* of this 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 generated by the scalar component of $\vec{L}$ along this

[18]The qualification given in footnote 15 from Chapter 1 applies here as well.

axis.

The technical reason for why Hamiltonian flows preserve the symplectic form (and thus count as state-space symmetries) need not concern us here. By contrast, the determinative relationship between smooth phase space functions and the Hamiltonian flows they generate provides crucial inspiration for the metaphysical account I will propose later, so bear with me while I explain this relationship in more detail.

5.1.2 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

[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.

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. In other words: the exterior derivative of a function is equivalent to this function up to an additive constant.

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. Intuitively: the role of the Riemannian metric is to take something that *doesn't* point in any direction (the exterior derivative of f) and turn it into something that *does* point in a direction (the gradient of f).

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? 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 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

[21]Since there is an infinity of distinct functions with the same exterior derivative—the functions that agree up to an arbitrary real constant—the flow of each gradient vector field has an infinity of generators. The same applies to Hamiltonian flows that will be introduced shortly.

field that everywhere points *along the level surfaces* of this function, as illustrated in figure (5.3). Vector fields of this sort are referred to as *Hamiltonian vector fields* and their flows

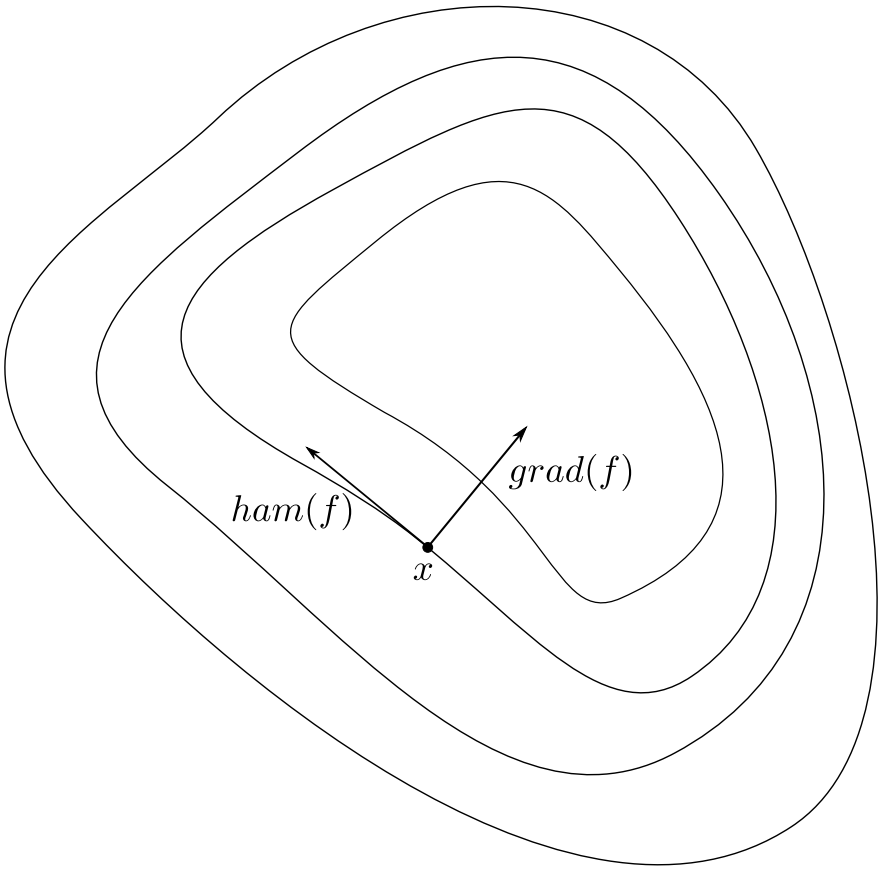


Figure 5.3: Illustration of the level surfaces (contour lines) of a function f on a two-dimensional manifold. $grad(f)$ and $ham(f)$ denote the values at x of the gradient and the Hamiltonian vector field of f, respectively.

.

as *Hamiltonian flows*. As a result, each phase space function f determines or *generates* a specific Hamiltonian flow: the flow of the Hamiltonian vector field of f. Moreover, since the Hamiltonian vector field of f points everywhere along the level surfaces of f, each flow line of the Hamiltonian flow generated by f is confined to some level surface of f. This is the reason for why smooth phase space functions are invariant under the Hamiltonian flows they generate.

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.[22] 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 Casimir invariant.

[22]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.

5.1.3 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 characterizes 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[23]

$$\frac{df}{d\lambda} = \{f, g\}, \tag{5.1}$$

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.[24]

For present purposes, the following intuitive picture of the mathematical content of (5.1) shall suffice. A perfectly standard way to represent the rate of change of the elevation of a

[23] 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.

[24] 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, $\{\{f, g\}, h\} + \{\{g, h\}, f\} + \{\{h, f\}, g\} = 0$.

piece of terrain along a path through the terrain is by means of the *derivative* of the elevation function in the direction of this path.[25]

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. 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 (5.4).

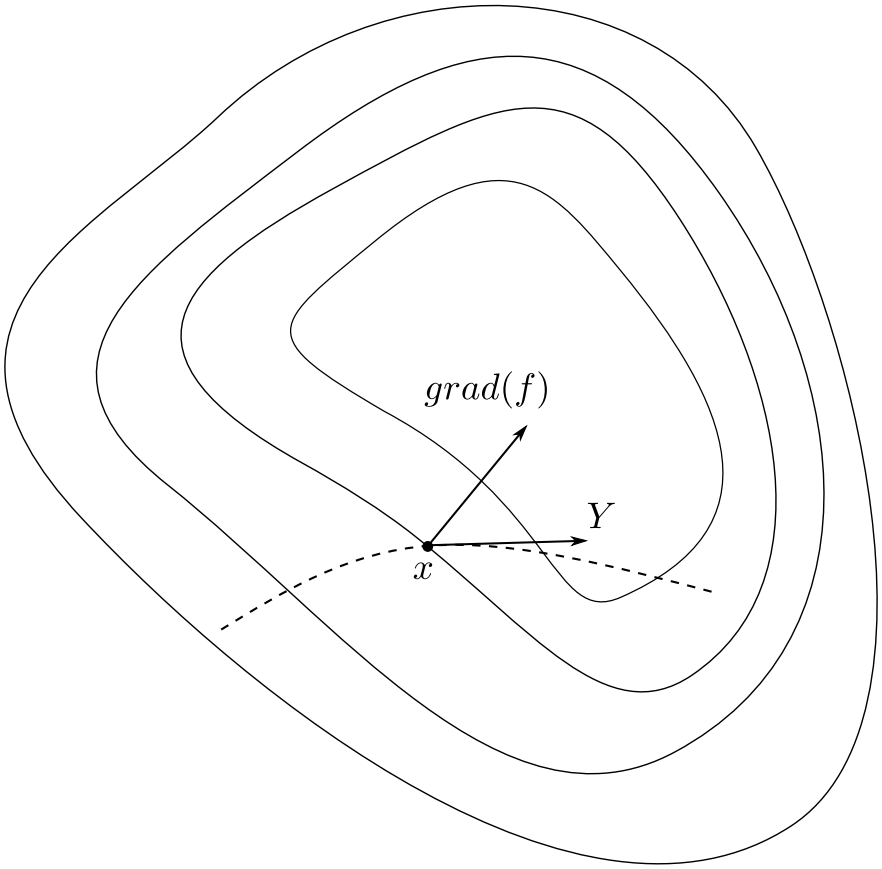


Figure 5.4: Illustration of a function f via its contour lines, showing a flow line (dotted) of the vector field Y and the value $grad(f)$ of the gradient of f at x.

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

Equation (5.1) 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 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 summarized as

$$\{L_i, L_j\} = \epsilon_{ijk} L_k, \tag{5.2}$$

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.[26]

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

[26] In mathematical jargon, equation (5.2) 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*.

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 axis i through x, its derivative along the Hamiltonian flow corresponding to rigid rotations about i vanishes: $\{f, L_i\} = 0$. 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. In other words: the quantities represented by SO(3)-constant functions are the *modal* SO(3)-invariants we encountered in the introduction to this chapter.

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.

5.1.4 Casimir Invariants

I follow the standard definition in physics texts.[27] 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.[28] 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,

[27](Penrose, 2004, pp. 553,568), (Fuchs and Schweigert, 2003, pp. 254).
[28]This is derived in (Goldstein et al., 2001, p. 418).

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 chapter, 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 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.

The goal of this section has been to introduce and explain the notion of a Casimir invariant, the mathematical representative of the target notion of a fine-grained invariant that we want to elucidate. The result is this: to be an SO(3)-Casimir is to be an SO(3)-constant function of the SO(3)-generators. The goal is now to characterize the notion of a fine-grained invariant by taking inspiration from its mathematical surrogate.

5.2 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 a fine-grained SO(3)-invariant from quantities that are merely modally invariant under SO(3). 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 whatever specific operation on the position and momentum facts at that state corresponds to taking 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, for M to be a fine-grained SO(3)-invariant is for it to be the case that 1) M is a modal SO(3)-invariant and 2) 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 (5.2) 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 determinate values of quantities represented by non-commuting operators.[29] This means

[29]Note that this claim does not depend on the controversial *eigenstate-eigenvalue link*. In other words: it

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

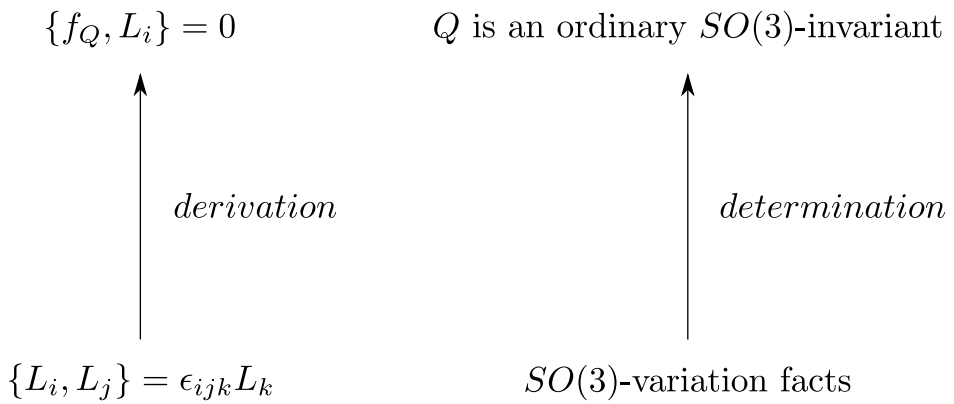


Figure 5.5: Illustration of my proposal, according to which the mathematical feature of the SO(3)-Casimir invariant function f_Q shown on the left is reflected in a similar feature, shown on the right, of the quantity Q represented by f_Q.

What, then, is the right account? The simple idea we just considered hypostatizes 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 the following proposal, for a property to be an SO(3) fine-grained invariant is for the modal SO(3)-invariance of this property to be *fully determined* by the

is not presupposed that (1) *every* physical quantity is represented by a Hermitian operator, nor that (2) if a quantity Q is represented by a Hermitian operator, the *only* way of determinately having Q is being in an eigenstate of Q.

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:

> **Fine-Grained SO(3)-Invariance.** For a quantity Q to be a *fine-grained SO(3)-invariant* is for modal SO(3)-invariance of Q to be fully determined by the SO(3)-variation facts.

To make this definition precise, we need to clarify the notions that occur in it: the notion of an SO(3)-variation facts and the notion of determination.

5.2.1 Modality, Variation, and Determination

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 with 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 (5.1.2) is suggestive of the following proposal. What is conveyed about a quantity Q by the fact that its mathematical surrogate f_Q is a generator of its associated phase-space flow is that the relevant 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, for some appropriate notion of determination.

To implement this proposal, we need to do three things. We need to understand the nature of the facts determined by certain facts about the quantities characterized by generators. We need to explain the nature of facts about those quantities that determine these facts. And we need to say more about the nature of the determination relation that obtains between the relevant facts.

What is the nature of the facts determined by certain facts about quantities characterized by 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 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 $\pm\phi$.[30] 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 M_1 are the 1-modal facts at every state.

Next, consider a physical quantity characterized by the generator of some phase-space flow. What is the nature of the facts about this quantity that are supposed to determine the relevant modal facts? Start by observing that the 1-modal facts at every state are characterized 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 (5.1.2), the Hamiltonian 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

[30] 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 $\frac{1}{360^\circ}|\phi|$ 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).

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, we need to say more about the nature of the determination relation. To begin with, determination in the relevant sense is a relation between *facts*: namely, the relation that obtains between (say) the 1-modal facts at some state and the level surface facts about M_1, together with the position and momentum facts at that state. Moreover, this relation is not identity: the conjunction of the relevant modal facts is clearly distinct from the conjunction of the corresponding level surface facts. It is also asymmetric: the relevant notion of determination is supposed to give metaphysical content to a piece of asymmetric mathematical language, namely the mathematical sense in which a smooth phase-space function is the *generator* of some specific phase-space flow. Thus, the target notion of determination is such that the 1-modal facts at some state are determined by, but do not determine, the level surface facts about M_1, together with the position and momentum facts at that state. The determination relation is also clearly not a causal relation: there is no obvious sense in which the 1-modal facts are caused by the level surface facts about M_1, together with the position and momentum facts.

Moreover, the relevant notion of determination is also distinct from necessitation and supervenience. I'll here focus on the case of supervenience with metaphysical necessity, according to which facts of type Y supervene on facts of type X iff it is metaphysically impossible for there to be a difference in the Y-facts without a difference in the X-facts. (The case against metaphysical necessitation is analogous.) The idea is as follows: 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 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. The relevant notion of determination is therefore more fine-grained than metaphysical necessitation and supervenience: although determination implies necessitation and supervenience, it is not tantamount to these notions.[31]

This feature of the determination relation is consistent with the following observation about its fineness of grain. 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 of the Hamiltonian flow of f, we have reason to deny that, if the level surface facts about some quantity Q (together with the position and momentum facts) determine the relevant modal facts, so do the level surface facts about every quantity state-space coextensive with Q. 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 facts about the *latter* determine the 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. The relevant notion of determination is more fine-grained than its mathematical surrogate.

Finally, the 'generating' idiom suggests that determination is an explanatory relation. For

[31]Note that this reasoning applies *a fortiori* to proposals which deploy narrower necessities, such as nomic necessity.

example: when we 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. This suggests that determination is an explanatory relation of a similarly productive (albeit non-causal) kind. According to this conception, when p determines 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*'.

So, the relevant notion of determination is that of an asymmetric non-causal explanatory relation between facts that implies (but is not implied by) metaphysical necessitation and supervenience.[32] These characteristics are compatible with various more specific metaphysical notions that have been discussed in the literature.[33] The project of this chapter, however, does not require a commitment to any particular such notion. For our purposes, the characterization of the fairly general features of the determination relation just given is sufficient.

This account of what is conveyed about a property by the fact that it is represented by a generator might seem abstract and unintuitive. But the general kind of explanatory move should be familiar from everyday modal claims. According to the account just presented, the fact that some proposition p is 1-possible is determined by 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 determined, and thus partially explained by, the fact

[32]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 determination claims. For example, the claim that the 1-modal facts are determined by 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 deterrmined by 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 the determination relation involve the relevant generating quantities, but not the definite descriptions by which these quantities can be picked out.

[33]The fineness of grain of the determination relation is suggestive of resources from hyperintensional metaphysics, such as metaphysical grounding (Fine, 2001, 2012; Rosen, 2010), essence (Fine, 1994a,b; Rosen, 2015), ontological dependence (Fine, 1995; Correia, 2008; Koslicki, 2012) or metaphysically impossible worlds (Nolan, 2013; Jago, 2014).

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.

We are now in a position to understand the definite descriptions that are expressed by Poisson brackets. 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 the state-space curves determined by the level surface facts about M_2'; and $\{L_1, L_2\} = L_3$ expresses the non-trivial identity of the satisfier of this definite description and M_3.[34]

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

Let me pause briefly to review what we have accomplished so far. Our goal is to understand what is conveyed about a physical quantity by the fact that it is represented by a Casimir invariant under some relevant group—say, SO(3). According to our strategy, the answer is that the modal SO(3)-invariance of that quantity is determined by the SO(3)-variation facts. The central result of this section is that we now have a precise account of the SO(3)-variation facts. The remaining task is to characterize the sense in which these facts *determine* the modal SO(3)-invariance of the relevant quantity. This is what I turn to now.

[34]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 determination claims are preserved under substitutions of co-denoting definite descriptions of the sort expressed by SO(3)-Poisson bracket equations.

5.2.2 Fine-Grained Invariance

There is reason to think that the relevant notion of determination relation is the same as the notion discussed in the previous section. Like before, determination is a relation between facts that is distinct from the identity relation: the conjunction of the SO(3)-variation facts is clearly not identical to the modal SO(3)-invariance of the magnitude of angular momentum. The relation is arguably asymmetric: the SO(3)-variation facts determine, but are not determined by, the fact that the magnitude of angular momentum is a modal SO(3)-invariant. It is also non-causal: there is no obvious sense in which facts of the former type *cause* facts of the latter type. Moreover, by familiar reasoning, determination is not necessitation or supervenience. 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 modal SO(3)-invariance of any vector quantity supervenes with metaphysical necessity on, and is metaphysically necessitated by, any fact whatsoever; but the modal SO(3)-invariance of this quantity is not determined by any fact whatsoever.

The relevant determination relation also has explanatory characteristics. 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 modal 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. Determination is an explanatory relation.

In sum, the operative notion of determination is, again, that of an asymmetric non-causal explanatory relation between facts that implies (but is not implied by) metaphysical

necessitation and supervenience.[35]

We are now in control of all notions involved in our account of fine-grained invariance, including the notion of an SO(3)-variation fact and the operative notion of determination. This means that our task is complete: we have sufficient understanding of the thesis that, for a quantity to be a fine-grained SO(3)-invariant is for the modal SO(3)-invariance of that quantity to be determined by the SO(3)-variation facts.

5.3 Ontic Structural Realism

Earlier, I noted that the phenomenon that motivates this chapter—assertions of a link between symmetries and properties encoded in the notion of a Casimir invariant—has inspired ontic structural realists to draw substantive metaphysical conclusions about the relative metaphysical priority of structures on the one hand and physical objects and their properties on the other. The core idea is that, if the essential properties of elementary particles can be exhaustively characterized in terms of mathematical objects that are defined in terms of symmetry groups—such as Casimir invariants—this suggests that symmetry structure is more fundamental than particles and their properties. To what extent does my account differ from ontic structuralist approaches?

To begin with, my ambitions in this chapter are more modest. The goal has been merely to explain what is conveyed about a physical quantity by the fact that it is mathematically characterized by a Casimir invariant, rather than to defend a claim about the relative fundamentality of such quantities and their bearers *vis-à-vis* symmetry structure. Moreover, the present investigation is not predicated on the goal of vindicating a structuralist metaphysics. The question as to whether the notion of a fine-grained invariant is consistent with ontic structuralism strikes me as open.

I take it that this is an advantage of the present investigation. One persistent criticism

[35] As before, for present purposes we can leave open which more specific metaphysical notion is best suited to play this role; see footnote 33 for some options.

of ontic structural realism has been that the operative notion of 'structure' has not been sufficiently clarified.[36] It therefore seems worth examining whether the metaphysical content of Casimir invariance can be identified without an antecedent commitment to ontic structural realism. This is what I have done in this chapter.

Another major difference between my account and ontic structural realist approaches concerns the status of mathematical objects in the specification of the relevant metaphysical claims. Structural realists have been forthright in their use of explicitly metaphysical notions to capture the relationship between symmetry structure and elementary particles: extant accounts deploy the language of necessitation,[37] supervenience,[38] determinables and determinates,[39] constitution,[40] ontological dependence,[41] fundamentality,[42] metaphysical grounding,[43] and essence.[44] However, the entities to which these notions are applied are often straightforwardly mathematical. A few examples: French (2014, p. 283) proposes that "the Poincaré group is a determinable". Roberts (2011, p. 50) argues for the claim that "the existing entities described by quantum theory are organized into a hierarchy, in which a particular symmetry group occupies the top, most fundamental position." And McKenzie (2014a, pp. 21-2) entertains the hypothesis that "the Poincaré group [...] forms part of [an elementary particle's] essence" which implies "the dependence of relativistic particles on the Poincaré group". What these claims have in common is that the subjects of the respective metaphysical attributes are mathematical entities introduced as part of the mathematical language used to state claims about physical reality.

As I explained in the introduction to this chapter, it is far from obvious what such claims amount to. At face value, the above structuralist theses have the implication that

[36] For example, (Psillos, 2001; Dorr, 2010; Wolff, 2011; Arenhart and Bueno, 2015); cf. (McKenzie, 2017).
[37] (Berenstain and Ladyman, 2012).
[38] (Ladyman et al., 2007; Wolff, 2011; Nounou, 2012).
[39] (French, 2014, Section 10.7-8).
[40] (Castellani, 1998a).
[41] (Wolff, 2011; McKenzie, 2014b).
[42] (Ladyman, 1998; Ladyman et al., 2007; Roberts, 2011; McKenzie, 2014b,a; French, 2014).
[43] (French, 2014, pp. 169,283).
[44] (McKenzie, 2014b).

symmetry groups—mathematical objects introduced as part of the mathematical apparatus of physics—are themselves among the constituents of physical reality. If this is not the claim these authors intend, we are owed a more transparent account of the ontological content of claims involving symmetry groups, an account which clearly and rigorously distinguishes between the underlying physical claims and the mathematical language used to express them.

One important goal of this chapter has been to respect this constraint, and the advantages of this approach should now be clear. We elucidated the physical content of what is perhaps one of the most important mathematical notions in elementary particle physics and thereby provided a transparent ontological basis for a symmetry-based metaphysics of modern physics. Moreover, our investigation demonstrates what we can achieve when we look beyond the mathematical apparatus in our ontological theorizing. For example, the fine-grained determinative relationship between the modal facts captured in terms of symmetry groups and the level surface facts of physical quantities is something we would likely have overlooked had we settled for the mathematical formalism of symmetry groups and their generators.

Finally, the account presented in this chapter creates conceptual space for a view which makes precise sense of the suggestion that the essential properties of particles are definable in terms of symmetries.[45] According to this view, for every essential property P of elementary particles, there is a continuous symmetry group G such that 'being a fine-grained G-invariant' defines P.[46] Applied to the classical-mechanical context of this chapter, the idea is that 'being a fine-grained SO(3)-invariant' defines the magnitude of angular momentum.[47] The notion of a fine-grained invariant developed in this chapter therefore promises a symmetry-based metaphysics of properties. Of course, the structuralist credentials such a view can only be fully assessed once it has been developed in detail. This is a task for another day.

[45]This suggestion has been made, for example, by Kantorovich (2003), Roberts (2011), McKenzie (2014a), French (2014), and Weinberg (1987, 1992). The original idea is based on a paper by Wigner (1939).

[46]To be sure: the sense of 'definition' operative in this view must be *metaphysical* rather than nominal: that is, the relevant definitions are definitions of *properties*, not of the predicates or concepts that refer to these properties. There are various accounts of metaphysical definition: see (Fine, 1994a; Rosen, 2015) for hyperintensional accounts and (Dorr, 2016) for an account in terms of higher-order identifications.

[47]This strategy has to be modified in the case of group actions which, unlike SO(3), have more than one fine-grained invariant. This issue should be taken up in future work.

Chapter 6

Conclusion

To conclude, I'd like to highlight a few general lessons to be drawn from this dissertation with regard to metaphysical theorizing about symmetries in particular and about physical theories in general.

First of all, in this work I hope to have demonstrated the power and utility of metaphysical theorizing about physics that puts symmetries front and center. For one thing, it is inspired by a fragment of mathematical physics that is at the core of virtually every successful fundamental theory developed over the last 100 years and thereby shows us the way toward a systematic and unified approach to the ontology of current and future physics. This is a significant achievement: *a priori*, there is little reason to think that theories as radically disparate as classical particle mechanics and loop quantum gravity would admit of a unified ontological framework. For another thing, symmetry-based theorizing allows us to provide hitherto unavailable explanations of a range of physical facts that lie at the heart of how modern physics explains the world around us—such as the discreteness of spin. Crucially, as chapter 3 demonstrated, the symmetry-based explanations of such facts do not amount to their mechancial reduction to abstract notions alien to physics; rather, they reveal the deep physical structure that underlies ordinary claims about the physical constituents of the world and their properties.

Second, ontological theorizing is often constrained by fairly tight principles that an ontology needs to satisfy for it to count as a genuine *candidate* for the ontology of a physical theory—for it to be plausible, intelligible, legitimate, or what have you. A good example if this is the constraint, oftentimes expressed by proponents of so-called 'primitive-ontology', that an acceptable fundamental ontology must involve local beables: roughly, physical entities confined to a bounded region of physical space-time. The investigation in this paper was not guided by constraints of this sort, and the advantages of this methodology should now be clear. An antecedent commitment (for example) to the Bohmian idea that local beables must be part of any reasonable fundamental ontology makes the entire project of giving precise content to symmetry fundamentalism a non-starter: after all, the symmetry-fundamentalist vision is one which demotes the material constituents of the world (such as particles) to ontologically second-rate status. As we have seen, symmetry fundamentalism points toward a powerful metaphysical framework for current and future physics alike—something we would likely have missed had we imposed more conventional constraints on ontological theorizing.

This brings me to the third remark. The metaphysical conclusions drawn in this dissertation were inspired by a segment of the mathematical formalism of physics that is regarded as central by practitioners themselves. But in so doing, we did not merely rest content in highlighting the fact *that* the relevant mathematical structures are physically or 'representationally' significant. This seemingly more metaphysically non-committal approach seems to be quite popular among philosophers of physics. Those sympathetic to this kind of approach are likely to think that I could have proceeded much more briskly throughout this dissertation: simply by pointing to the relevant mathematical structures—symmetries, symmetry groups, and Casimir invariants—and announced that it is the physical entities to which those mathematical notions correspond, *whatever they may be*, that are the subjects of the relevant claims involving fundamentality or ontological priority. It is an implicit premise of this approach that identifying which mathematical items are representationally significant is all that can be said by way of specifying an ontology: nothing further is to be gained by going

'beyond' the mathematical formulations in the way exhibited in this dissertation.

The approach I have taken here is more ambitious. We have characterized ontological items in physical terms; that is, in a way that makes transparent their distinctness from their mathematical representations and that vindicates the appropriateness of those representations for characterizing these items.[1] The advantages of this more ambitious approach should be clear. First of all, we discovered that the relevant physical structures across various physical theories are much simpler than their representation in terms of the complex formalism of symmetries and symmetry groups might suggest. For example, the physical structure of the fundamental state space in classical and quantum versions of symmetry fundamentalism can be captured just in terms of two symmetry relations—the respective almost-sameness relations—plus a sparse range of other primitives. Similarly, we saw that the physical structure on a state space specified by a symmetry group action can be captured just in terms of a single four-place relation; and the physical content of a Casimir invariant under some symmetry can be captured in terms of a determination relation that facts about modal invariances under that symmetry bear to facts about the variation of specific physical properties under that symmetry. This considerable simplicity in the fundamental physical structure characterized by mathematical symmetries is something that we would likely have missed had we not gone beyond the mathematical formalism in our ontological theorizing.

Second of all, the standard representationalist strategy breaks down in the face of metaphysical theses as radical and seemingly abstract as symmetry fundamentalism: in such cases, this approach simply cannot provide us with anything close to the metaphysical insight we're looking for. There are settings in which we are sufficiently in control of the metaphysical background assumptions and the supporting intuitions to securely follow the representationalist strategy. I take it that most of us would have a fairly clear picture of what is being asserted

[1]This approach is inspired by nominalistic constraints on ontological theorizing, according to which physical ontology must be specified without invoking mathematical entities at all; cf. (Field, 1980; Arntzenius and Dorr, 2012; Chen, 2018). However, limitations of space prevent me from fully complying with nominalistic strictures in this dissertation. Instead, I have provided a nominalistic characterization of merely the relevant parts of the ontology: those concerned with the items corresponding to mathematical symmetry structures.

by someone who tells us that the fundamental arena in which the history of the world unfolds is characterized by a three-dimensional manifold with a Euclidean metric. The claim that symmetries are fundamental and that every material object is ontologically derivative of them is quite different. Being told that a range of symmetry transformations are representationally significant and 'correspond to' physical items does not give us anything like a sufficiently clear picture of what the world would have to be like for that claim to be true. Indeed, this person would run the risk of being quite profoundly misunderstood: to some, it may sound like we're proposing a kind of pythagoreanism, according to which physical reality consists (at least in part) of mathematical objects. Specifying the relevant ontological primitives in purely physical terms—as I did in throughout this dissertation—is a particularly powerful way to avoid problems of this kind. The present work therefore highlights the importance of ontological theorizing that looks beyond the mathematical formulations of physical theories in order to discern the physical structures characterized by them.

Finally, our investigation demonstrates that taking physics and its mathematical formalism as the starting point for metaphysical theorizing isn't a way of making metaphysics or its methods redundant, as some would have you believe. Quite the contrary: doing so showcases the power of metaphysics as a tool for exploring the full range of ontological options for fundamental physics.

Bibliography

Albert, D. (1996). Elementary Quantum Metaphysics. In Cushing, J. T., Fine, A., and Goldstein, S., editors, *Bohmian Mechanics and Quantum Theory: An Appraisal*, pages 277–284. Kluwer Academic Publishers.

Allori, V. (2013). Primitive Ontology and the Structure of Physical Theories. In Ney, A. and Albert, D. Z., editors, *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*, pages 58–75. Oxford University Press.

Arenhart, J. R. B. and Bueno, O. (2015). Structural realism and the nature of structure. *European Journal for Philosophy of Science*, 5(1):111–39.

Arntzenius, F. and Dorr, C. (2012). Calculus as geometry. In Arntzenius, F., editor, *Space, Time, and Stuff*, chapter 8, pages 213–78. Oxford: Oxford University Press.

Ashtekar, A. and Schilling, T. A. (1999). Geometrical Formulations of Quantum Mechanics. In *On Einstein's Path*. Springer. arXiv:gr-qc/9706069v.

Audi, P. (2012). Grounding: Toward a Theory of the In-Virtue-Of Relation. *Journal of Philosophy*, 109(12):685–711.

Bacon, A. (2019). Is reality fundamentally qualitative? *Philosophical Studies*, 176(1):259–95.

Ballentine, L. E. (2000). *Quantum Mechanics: A Modern Development*. London: World Scientific.

Belot, G. (2012). Quantum states for primitive ontologists: A case study. *European Journal for Philosophy of Science*, 2(1):67–83.

Berenstain, N. and Ladyman, J. (2012). Ontic Structural Realism and Modality. In Landry, E. M. and Rickles, D. P., editors, *Structural Realism*, pages 149–168. Springer.

Brading, K. and Castellani, E. (2003). *Symmetries in Physics – Philosophical Reflections*. Cambridge: Cambridge University Press.

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.

Brading, K., Castellani, E., and Teh, N. (2017). Symmetry and Symmetry Breaking. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2017 edition.

Brown, H. R. and Wallace, D. (2005). Solving the Measurement Problem: De Broglie-Bohm Loses Out to Everett. *Foundations of Physics*, 35(4):517–540.

Bueno, O. (2011). Structural Empiricism, Again. In Bokulich, A. and Bokulich, P., editors, *Scientific Structuralism*, volume 281 of *Boston Studies in the Philosophy of Science*, pages 81–103. Springer.

Castellani, E. (1998a). Galilean Particles: An Example of Constitution of Objects. In *Interpreting Bodies*, pages 181–194. Princeton University Press.

Castellani, E. (1998b). *Interpreting Bodies*. Princeton: Princeton University Press.

Chen, E. K. (2017). Our Fundamental Physical Space: An Essay on the Metaphysics of the Wave Function. *Journal of Philosophy*, 114(7):333–365.

Chen, E. K. (2018). The Intrinsic Structure of Quantum Mechanics. *PhilSci Archive*. Article ID: 15140.

Chen, E. K. (2019). Realism about the wave function. *Philosophy Compass*. `https://doi.org/10.1111/phc3.12611`.

Correia, F. (2008). Ontological Dependence. *Philosophy Compass*, 3(5):1013–1032.

Dasgupta, S. (2009). Individuals: an essay in revisionary metaphysics. *Philosophical Studies*, 145(1):35–67.

Dasgupta, S. (2014). The Possibility of Physicalism. *Journal of Philosophy*, 111(9-10):557–592.

Dorr, C. (2010). Review of Every Thing Must Go: Metaphysics Naturalized. *Notre Dame Philosophical Reviews*. `https://ndpr.nd.edu/news/every-thing-must-go-metaphysics-naturalized/`, accessed 22 April 2020.

Dorr, C. (2016). To Be F Is To Be G. *Philosophical Perspectives*, 30(1):39–134.

Dorr, C. (2018). A Way of Getting Rid of Things: Higher-order Languages, Priorian Nominalism, and Nihilism. `http://www.nyu.edu/projects/dorr/handouts/Nihilism.pdf`, accessed 22 April 2020.

Earman, J. (2003). Rough guide to spontaneous symmetry breaking. In Brading, K. and Castellani, E., editors, *Symmetries in Physics–Philosophical Reflections*, pages 335–46. Cambridge: Cambridge University Press.

Field, H. (1980). *Science Without Numbers*. Princeton: Princeton University Press.

Fine, K. (1994a). Essence and Modality. *Philosophical Perspectives*, 8:1–16.

Fine, K. (1994b). Senses of Essence. In Sinnott-Armstrong, W., Raffman, D., and Asher, N., editors, *Modality, Morality, and Belief – Essays in Honor of Ruth Barcan Marcus*, pages 53–73. Cambridge University Press.

Fine, K. (1995). Ontological Dependence. *Proceedings of the Aristotelian Society*, 95:269–290.

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.

Folland, G. (2015). *A Course in Abstract Harmonic Analysis*. Boca Raton: CRC Press, 2nd edition.

Forrest, P. (1988). *Quantum Metaphysics*. Blackwell.

French, S. (2010). The interdependence of structure, objects and dependence. *Synthese*, 175(1):89–109.

French, S. (2013). Whither Wave Function Realism? In Ney, A. and Albert, D. Z., editors, *The Wave Function – Essays on the Metaphysics of Quantum Mechanics*, pages 76–90. Oxford University Press.

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

French, S. and Ladyman, J. (2003). Remodelling Structural Realism: Quantum Physics and the Metaphysics of Structure. *Synthese*, 136(1):31–56.

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.

Hawthorne, J. (2010). A Metaphysician Looks at the Everett Interpretation. In Saunders, S., Barrett, J., Kent, A., and Wallace, D., editors, *Many Worlds? Everett, Quantum Theory & Reality*. Oxford University Press.

Heisenberg, W. (1973). Development of Concepts in the History of Quantum Theory. In Mehra, J., editor, *The Physicist's Conception of Nature*, pages 264–75. Dordrecht: Springer.

Heisenberg, W. (1976). The nature of elementary particles. *Physics Today*, 29(3):32–9.

Hubert, M. and Romano, D. (2018). The Wave-Function as a Multi-Field. *European Journal for Philosophy of Science*, 8(3):521–537.

Ismael, J. and Schaffer, J. (2016). Quantum holism: nonseparability as common ground. *Synthese*, Special Issue: The Metaphysics of Entanglement. `https://doi.org/10.1007/s11229-016-1201-2`.

Jago, M. (2014). *The Impossible: An Essay on Hyperintensionality*. Oxford University Press.

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):651–75.

Kim, Y. S. and Noz, M. (2012). *Theory and Applications of the Poincaré Group*. D. Reidel Publishing Company.

Klein, F. (1893). Vergleichende Betrachtungen über neuere geometrische Forschungen. *Mathematische Annalen*, 43(1):63–100. Revised version of the booklet by the same title published by A. Duchert, Erlangen, in 1872.

Klein, F. (1910). Über die geometrischen Grundlagen der Lorentz-Gruppe. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 19:281–300. `https://eudml.org/doc/145243`, accessed 22 April 2020.

Kochen, S. (1996). Construction of Quantum Mechanics Via Commutative Operations. In Clifton, R., editor, *Perspectives on Quantum Reality*, volume 57 of *The University of Western Ontario Series in the Philosophy of Science*, pages 237–43. Dordrecht: Springer.

Koslicki, K. (2012). Varieties of ontological dependence. In *Metaphysical Grounding*, chapter 7, pages 186–213. Cambridge University Press.

Ladyman, J. (1998). What is structural realism? *Studies in the History and Philosophy of Science Part A*, 29(3):409–24.

Ladyman, J. (2010). Reply to Hawthorne: Physics before Metaphysics. In Saunders, S., Barrett, J., Kent, A., and Wallace, D., editors, *Many Worlds? Everett, Quantum Theory & Reality*, pages 154–160. Oxford University Press.

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

Loewer, B. (1996). Humean Supervenience. *Philosophical Topics*, 24(1):101–127.

Martin, C. A. (2003). On continuous symmetries and the foundations of modern physics. In Brading, K. and Castellani, E., editors, *Symmetries in Physics–Philosophical Reflections*, pages 29–61. Cambridge University Press.

Maudlin, T. (2010). Can the World be Only Wavefunction? In Saunders, S., Barrett, J., Kent, A., and Wallace, D., editors, *Many Worlds? Everett, Quantum Theory, and Reality*, chapter 4, pages 121–143. Oxford University Press.

Maudlin, T. (2013). The Nature of the Quantum State. In Ney, A. and Albert, D., editors, *The Wave Function—Essays on the Metaphysics of Quantum Mechanics*, pages 126–153. Oxford University Press.

Maudlin, T. (2018). Ontological Clarity via Canonical Presentation: Electromagnetism and the Aharonov-Bohm Effect. *Entropy*, 20(6). 465.

McKenzie, K. (2014a). On the Fundamentality of Symmetries. *Philosophy of Science*, 81(5):1090–102.

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

McKenzie, K. (2017). Relativities of fundamentality. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 59:89–99.

Moretti, V. (2013). *Spectral Theory and Quantum Mechanics: With an Introduction to the Algebraic Formulation*. Milan: Springer.

Myrvold, W. (2018). Philosophical Issues in Quantum Theory. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, fall 2018 edition.

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.

Ney, A. (2012). The Status of Our Ordinary Three Dimensions in a Quantum Universe. *Noûs*, 46(3):525–560.

Ney, A. (2013). Introduction. In Ney, A. and Albert, D. Z., editors, *The Wave Function – Essays on the Metaphysics of Quantum Mechanics*, pages 1–51. Oxford University Press.

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

Nolan, D. P. (2013). Impossible Worlds. *Philosophy Compass*, 8(4):360–372.

North, J. (2009). The "Structure" of Physics: A Case Study. *The Journal of Philosophy*, 106(2):57–88.

North, J. (2013). The Structure of a Quantum World. In Ney, A. and Albert, D. Z., editors, *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*, pages 184–202. Oxford University Press Oxford.

Nounou, A. M. (2012). Kinds of Objects and Varieties of Properties. In Landry, E. M. and Rickles, D. P., editors, *Structural Realism*, pages 117–133. Springer.

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

Psillos, S. (2001). Is Structural Realism Possible? *Philosophy of Science*, 68(3):S13–S24.

Roberts, B. W. (2011). Group Structural Realism. *The British Journal for the Philosophy of Science*, 62(1):47–69.

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.

Rosen, G. (2015). Real Definition. *Analytical Philosophy*, 56(3):189–209.

Rovelli, C. (2011). Zakopane lectures on loop gravity. *arXiv:1102.3660*. `https://arxiv.org/abs/1102.3660`, accessed 22 April 2020.

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

Saunders, S. (2010). Many Worlds? An Introduction. In Saunders, S., Barrett, J., Kent, A., and Wallace, D., editors, *Many Worlds? Everett, Quantum Theory, and Reality*, pages 1–49. Oxford University Press.

Schaffer, J. (2009). On What Grounds What. In Manley, D., Chalmers, D. J., and Wasserman, R., editors, *Metametaphysics: New Essays on the Foundations of Ontology*, pages 347–383. Oxford University Press.

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

Strocchi, F. (2005). *Symmetry Breaking*. Berlin: Springer.

Swanson, N. (2018). How to Be a Relativistic Spacetime State Realist. *The British Journal for the Philosophy of Science*. axy041.

Wallace, D. (2008). Philosophy of Quantum Mechanics. In Rickles, D., editor, *The Ashgate Companion to Contemporary Philosophy of Physics*, pages 16–98. Ashgate.

Wallace, D. (2012). *The Emergent Multiverse: Quantum Theory According to the Everett Interpretation*. Oxford University Press.

Wallace, D. (2013). A Prolegomenon to the Ontology of the Everett Interpretation. *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*, Ney, Alyssa and Albert, David Z.:203–222.

Wallace, D. (2017). Against Wave-Function Realism. *(manuscript)*.

Wallace, D. (2019). Who's afraid of coordinate systems? An essay on representation of spacetime structure. *Studies in the History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 67:125–36.

Wallace, D. and Timpson, C. G. (2010). Quantum Mechanics on Spacetime I: Spacetime State Realism. *British Journal for the Philosophy of Science*, 61(4):697–727.

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: Cambridge University Press.

Weinberg, S. (1992). *Dreams of a Final Theory*. New York: Pantheon Books.

Weinberg, S. (1995). *The Quantum Theory of Fields Vol. I*. Cambridge: Cambridge University Press.

Wigner, E. P. (1931). *Gruppentheorie und ihre Anwendung auf die Quantenmechanik der Atomspektren*. Braunschweig: Vieweg.

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

Wigner, E. P. (1968). Symmetry principles in old and new physics. *Bulletin of the American Mathematical Society*, 74(5):793–815.

Williamson, T. (2016). Modal science. *Canadian Journal of Philosophy*, 46(4-5):453–92.

Wolff, J. (2011). Do Objects Depend on Structures? *The British Journal for the Philosophy of Science*, 63(3):607–25.