

Varieties of Wave Function Realism, or, WTF is WFR?

Wayne C. Myrvold
Department of Philosophy
The University of Western Ontario
wmyrvold@uwo.ca

April 30, 2025

Contents

1	Introduction: How many varieties of WFR?	3
2	Varieties of $3N$-dimensional space	9
3	Preliminaries: Classical fields on space-time	13
4	The plurality of wave function representations of a quantum state	18
5	The cornucopia of wave function representations of quantum states	20
6	Dynamics	25
7	The fundamental arena as a differentiable manifold	26
8	Wave function realisms	31
9	Conclusion	32
10	Acknowledgments	33
11	Appendix: the plurality of wave-function representations	35
11.1	Hilbert spaces from function spaces	35
11.2	Unitarily equivalent representations	36

1 Introduction: How many varieties of WFR?

There has been considerable discussion in the philosophical literature of the past decade or so of a view that has come to be known as “wave function realism,” which I will abbreviate as *WFR*. The basic claim of this view is that quantum theory gives us motivation to think that quantum wave functions should be thought of as fields on a space of very high (or perhaps infinite) dimension, and that this space is in some important sense more fundamental than familiar three-dimensional space or four-dimensional spacetime. Note that this is much stronger than the mere claim that quantum states represent *something* physically real, a claim that I myself have defended (Myrvold 2020a, 2020b).

The view is inspired by remarks made by J. S. Bell, in lectures and conversations though not in print, to the effect that, because a wave function for an N -particle system is a function on a $3N$ -dimensional space, “somehow our feeling that we live in $3 + 1$ dimensions is an illusion.”¹ It entered the philosophical literature in David Albert’s “Elementary Quantum Metaphysics” (1996),² and since then has gained some currency. According to Alyssa Ney, “it has had something of an outsized acceptance in metaphysical circles” (2021b, xi, fn. 3) although, within philosophy of physics, its advocates “can be counted on a single hand” (xii).

Albert’s original article gave the impression that the conclusions he drew followed straightforwardly from taking a realist stance towards quantum mechanics; that is, it was suggested that, once the pernicious *Copenhagener Geist* has been exorcised, and we allow ourselves to ask what quantum mechanics is telling us about the physical world, it follows that the impression we have of living in a three-dimensional space, or a four-dimensional spacetime, “is somehow flatly illusory” (Albert, 1996, 277).³ We find a similar conclusion in Loewer (1996), where the position is explicitly motivated by the enterprise of adapting D. K. Lewis’ Humean supervenience to the quantum context. Criticisms were advanced by Monton (2002) and Lewis (2004). A significant advance in the discussions was the publication in 2013 of a volume dedicated to discussions, pro and con, of the view (Ney and Albert, 2013). More recently, WFR has been the subject of a book-length defense

¹See Bell (1990), remarks beginning at 1:25:15.

²Reprinted, with modifications, as Albert (2013).

³In Albert’s subsequent expositions (2013; 2015; 2019), talk of *illusion* has been dropped. See section 7, below.

by Alyssa Ney (2021b).

In my article, “What is a Wave Function?” (Myrvold, 2015), I attempted to undermine the motivation for a view of this sort by invoking a principle to the effect that ontological conclusions drawn from quantum mechanics should be compatible with the uncontroversial fact that quantum mechanics is not a fundamental theory, but a non-relativistic approximation to a more fundamental theory, quantum field theory. This means that whatever it is in the physical world that quantum-mechanical wave functions represent must be approximations to something that is defined in terms of the underlying quantum field theory. If one accepts that principle, then our thinking about the ontological status of wave functions should be informed by our understanding of how quantum mechanics is obtained, in a non-relativistic regime, from a quantum field theory. This is relevant to the issue, because quantum-mechanical wave functions are defined in terms of field operators, $\hat{\phi}(x)$, defined on ordinary 4-dimensional spacetime. The N -particle wave-functions that emerge in a nonrelativistic regime assign numbers to N -tuples of points in four-dimensional spacetime, not to points in some high-dimensional space whose relation to ordinary spacetime is obscure.

In this chapter, my aim is not to criticize wave function realism, but to get clearer about what it is. In particular, I want to ask whether the phrase “wave function realism” denotes a monolithic position, or whether there are multiple varieties of wave function realism. One purpose is to point out that, even if all current proponents of WFR were to agree on all essential points, there is still at least a potential plurality of wave function realisms—in principle, at least, there is more than one variety of wave function realism. Another purpose is to highlight some questions about the nature of the supposed fundamental space, and the wave functions that are defined on it, that need to be answered before we know what the position is supposed to amount to. Until these have been answered, then, as far as explicitly articulated versions of WFR go, we have, not more than one version of WFR, but less than one.

I will focus on two main axes of distinction. One has to do with the nature of the project. Is it what I will call *Interpretive* project, one of accepting standard quantum theory pretty much as we have it, and exploring its implications for ontology? Or is it a *Constructive* project, which finds standard quantum theory wanting in crucial aspects, and seeks to construct a new theory that will satisfy some set of metaphysical constraints? The other has to do with how radical the claims are that are made about the nature of space-time. Does the fundamental space on which the wave functions of WFR are

defined have intrinsic structure corresponding to the low-dimensional space-time structure? One flavour of WFR, an *Extremely Mild* one, would have it that the fundamental space has, as intrinsic structure, a representation of the symmetry group of the low-dimensional space, and that the dynamics of the wave function are, as a matter of physical law, required to respect these symmetries. On such a view, it would be hard to see in what sense the low-dimensional spatial structure is not fundamental; any sense in which the low-dimensional spatial structure is non-fundamental would be at best a highly attenuated one. A more radical view, a *Spicy* flavour of WFR, would have it that the fundamental space has no structure at all corresponding to the low-dimensional spatial structure, and that no conditions are placed on the dynamics of wave functions that amount to a condition of respecting the symmetries of some low-dimensional space. According to a Spicy version of WFR, if it turns out that a wave function evolving on a $3N$ -dimensional fundamental space can be simply represented as a function of N -tuples of points in a 3-dimensional Euclidean space, evolving according to a dynamical law that depends only on distances between those N points in the 3-dimensional space, this is a contingent fact, and there are worlds that are possible, according to the theory, in which the wave function has nothing like these features. On this view, wave functions in general don't have any relation to any low-dimensional space or spacetime.

Much of the literature on WFR suggests that proponents of WFR intend it to be Interpretive and Spicy. As will be explained in sections 4 and 5, there are serious difficulties faced by a position of this sort, and it is at least *prima facie* an untenable combination. Standard quantum theory makes use of the low-dimensional spatial or spacetime structure to formulate wave functions. If this structure is stripped away, and we have to do with functions defined on a high-dimensional space that can be specified without reference to the low-dimensional structure, these functions will have to be *unlike* the wave functions of standard quantum mechanics in crucial ways. That is, a Spicy version of WFR cannot be merely Interpretive.

There is another way in which WFR, as expounded by its proponents, departs from standard quantum theory. There is often talk of *the* wave function, in the singular, that represents a given quantum state. It's not always clear whether this is meant to be taken literally. If it is, this is a radical departure from standard quantum theory, on which there is no such thing as *the* wave function that represents a quantum state. For any quantum state there is, rather, a veritable cornucopia of wave functions that can be

used to represent the state; we will explore this cornucopia in section 5, below.

A quantum state, whatever else it might do, encodes probabilities of outcomes of experiments. One has a choice of what mathematical apparatus to employ to do this, and students of quantum mechanics learn a variety of different ways to represent a quantum state, all agreeing on the probabilities of outcomes of all possible experiments. Standard quantum theory takes all of these to be *physically* equivalent: if two representations of a quantum state agree on all probabilities of outcomes of any experiments that could in principle be performed, they are representations of the same quantum state, and any features that the two representations do not have in common are thought not to have physical significance. A familiar example of this, often mentioned in textbooks, is that two wave functions that differ only by a multiplicative constant represent the same quantum state; changing a wave function by multiplying it a constant does not correspond to a change in what is represented by the wave function. Also familiar from textbooks is freedom of choice of basis; the same quantum state could, for example, be represented by a position-space wave function, or a momentum-space wave function.

One could, of course, construct a theory on which two wave functions representing the same quantum state are not physically equivalent. The project of doing so is clearly a Constructivist project, not one that is merely Interpretive.

It does seem to be essential to the WFR project that it have radical implications for our conception of spacetime, that is, that it be Spicy. As I will argue, it will then also have to be Constructive. This isn't an objection. Other avenues of approach to the so-called measurement problem of quantum mechanics, including hidden-variables approaches and dynamical collapse theories, are straightforwardly Constructive projects. The point is, rather, that, if WFR is also a Constructive project, it would be better to be explicit about this, and to be clear about what the theory to be constructed is. This would involve answering crucial questions about the structure of the fundamental space, what it is that the wave functions of WFR assign to points in that space, and what the dynamical laws are that govern their evolution. Proponents of WFR have often been somewhat vague on these points.

There's another, very substantial lacuna in the literature on WFR. As David Wallace (2021) has emphasized, expositions of WFR tend to be framed

in terms of a specific kind of quantum theory, the nonrelativistic quantum theory of a finite number of spinless particles. This is evidenced by the fact that it is said that a wave-function assigns a *number* to points in configuration space. In standard quantum mechanics, what a configuration-space wave function assigns to a point in configuration is a vector in the Hilbert space of total spin for the system. This Hilbert space has dimension equal to the product of the dimensions of the spin spaces of the individual particles. So, for example, for a system of N spin- $\frac{1}{2}$ particles, each particle's spin-space is a two-dimensional Hilbert space, and so the total spin-space of the system is a Hilbert space of dimension 2^N . This grows exponentially with N , so, if N is a mind-bogglingly large number, 2^N is exponentially more so. Alternatively, one could choose a basis for the spin space, and represent a vector in the spin-space by an ordered list of 2^N numbers, yielding 2^N wave functions. Any version of WFR for nonrelativistic quantum mechanics will have to specify what it is that is assigned to points in configuration space. Is there a single wave function that assigns, to each point in configuration space, a vector in a Hilbert space of huge dimension, or is there a huge number of complex-valued wave functions? Or is it something else, that we haven't thought of?

If what the wave functions of WFR assign to points in configuration space are vectors in the total spin-space of the system, this is a version of WFR that is at least Somewhat Mild, because the Hilbert space representation of spin is tied up with directions in three-dimensional space. One way to see this is to consider the spin state of a pair of spin-1/2 particles that is often used to illustrate entanglement, that is, the singlet state. This is a state that is not a state of definite spin, in any direction, for either of the component particles, but is a state of definite total spin for the pair of particles, in any direction. That is, if spin is measured *in the same direction* on both of the particles, the two results have to add up to zero. To even say that—that is, to even assert that the spin state of the pair of particles is a singlet state—we need to know how to make sense of *same direction* in three-dimensional space for the two spins. (If we measure spin in two different directions, the results are not guaranteed to add to zero in the singlet state). Furthermore, in standard quantum mechanics, operators corresponding to spin in three mutually perpendicular directions are required to satisfy commutation relations that are based on the commutation relations for components of angular momentum in three mutually perpendicular directions in three-dimensional space.⁴

⁴I am grateful to James Ladyman for stressing this latter point.

Of course, non-relativistic quantum mechanics, useful as it is, is not adequate in the relativistic regime. There are a number of arguments in the literature to the effect that a relativistic quantum theory will have to be a *quantum field theory*, that is, a quantum theory of a system of infinitely many degrees of freedom. What WFR is supposed to say about quantum field theories remains largely unexplored. Proponents tend to say that the fundamental space will be infinite-dimensional. But we have yet to see a concrete proposal as to the structure of this infinite-dimensional fundamental space. One possibility, for bosonic fields, is that the space on which the wave-function is to be defined is to be a space of possible field configurations. But a straightforward reading of that is that it is a space of possible assignments of field values to points in three-dimensional space or four-dimensional space-time. If that's the case, it appears that the structure of the low-dimensional space is baked into the characterization of the fundamental space, and once again we have a Mild version of WFR (see §4.3 of Myrvold 2015 for further discussion).

At any rate, proponents of WFR owe us an account of the fundamental space for a quantum field theory, its relation to four-dimensional space time, and its relation to the fundamental space of nonrelativistic quantum mechanics. It won't do to say that we have one fundamental space for nonrelativistic quantum mechanics, and another for quantum field theories, with no relation between the two, because the systems treated of in nonrelativistic quantum mechanics are also treated of in quantum field theories. When one calculates, in quantum field theory, the magnetic moment of the electron to astonishing precision, this is the magnetic moment of the same electron that is dealt with in elementary quantum-mechanical treatments of the hydrogen atom. What one is doing, in quantum field theories, is treating of relativistic phenomena that can be disregarded at low energies; one is not investigating some other world distinct from that dealt with in nonrelativistic quantum mechanics.

These are matters that will have to be sorted, before we have even one candidate version of WFR on the table. In this chapter, I will focus on what seems to be the easiest case for WFR, the case that is invariably used to motivate the view: the nonrelativistic theory of a finite number of spinless particles. Readers should bear in mind that answering the questions posed for this simplest case, though necessary, is not sufficient for the project of formulating an empirically adequate version of WFR.

Some of what has been said in this introductory section will have been unclear to some readers. I ask those readers to be patient, and read on. In

section 2 the potential ambiguities in the structure of the fundamental space will be spelled out in some detail. In section 3, as a prelude to a discussion of the dependence of the formulation of wave functions on the structure of the spacetime within which the theory is formulated, we will first review the more familiar territory of how this works with classical fields. As we will see in section 4, much of what is said in connection with classical fields holds also for quantum wave functions. This will start to give us a sense of the variety of wave functions that can all represent the same quantum state. We will outline the full rich abundance of wave function representations of quantum states in section 5.

2 Varieties of $3N$ -dimensional space

It is often said that the space on which the wave functions of WFR are defined is, at least for the case of N distinguishable particles, a $3N$ -dimensional space with the structure of a classical configuration space. There is potential ambiguity here, as there are different $3N$ -dimensional spaces that could be meant.

Let's start with the notion of a configuration of N bodies, treated classically and nonrelativistically. We might, for example, be engaged in an analysis, within classical physics, of the solar system. Suppose that N bodies—say, the sun and the major planets, and some finite number of other solar system objects—have been selected as the ones whose motions we are interested in, and, at the level of analysis employed, internal structure of the bodies under consideration may be ignored, so that it is enough to keep track of the motions of the centers of mass of all of the bodies considered.

We model the instantaneous positions of these bodies as an ordered N -tuple of points in three-dimensional Euclidean space, which we will call \mathbb{E}^3 . Ordered, because the bodies under consideration have distinguishing characteristics; swapping the positions of Saturn and Jupiter yields a distinct configuration. The space \mathbb{E}^3 has a lot of structure. Its structure permits distinction between smooth curves in the space and jagged ones, and, among the smooth curves, distinction between straight and curved paths. There is a notion of rigid displacement, which underwrites judgments of congruence of regions of this space; two regions are congruent if they can be made to coincide via a rigid displacement. Given any two pairs of points (a, b) and (c, d) , there is a matter of fact about whether or not the line segment joining

a and b is congruent to the line segment joining c and d . If they are not congruent, one of them is shorter than the other, meaning that the shorter is congruent to a proper part of the longer. It thus makes sense, given the structure of the space, to say that one line segment is twice as long as another, or n times as long, and this gives us means of saying, for any two line segments and any natural numbers n, m , whether or not n times the one is longer, shorter, or the same length as m times the other, and that gives us the means of associating, with any two line segments, a unique real number that is the ratio of their lengths. If, we now choose some line segment to serve us a unit length, we have, for any pair of points, p, q , a unique real number that is the ratio of the length of the line segment joining p and q to that of our unit length. Rigid displacements include rotations about any line, and this permits to assign an angle of intersection to any pair of lines that intersect. Two lines are perpendicular if their angle of intersection is one-quarter of a full rotation.

If we choose a point as an origin, a triple of mutually perpendicular directions in space, and a unit of length, we have a coordinatization of \mathbb{E}^3 that indicates points in the space via triples of real numbers, (x, y, z) . Because of the structure of our space, our choice of coordinates only requires three things: a choice of origin, a choice of three mutually perpendicular directions from that point, and a choice of unit. If we had a space with less intrinsic structure, we might still be able to coordinatize it, but this would involve a greater number of arbitrary choices.

We will call the space $(\mathbb{E}^3)^N$, which is the set of all ordered N -tuples of points in \mathbb{E}^3 , the *Configuration Space* of the system.⁵ Given a coordinatization of 3D Euclidean space \mathbb{E}^3 , we get an induced coordinatization of the Configuration Space $(\mathbb{E}^3)^N$, on which a point in $(\mathbb{E}^3)^N$ is represented by an ordered N -tuple of triples of real numbers, that is, something of the form $((x_1, y_1, z_1), \dots, (x_N, y_N, z_N))$.

The space $(\mathbb{E}^3)^N$, being the set of ordered N -tuples of points in \mathbb{E}^3 , inherits a lot of structure from it. An element p of $(\mathbb{E}^3)^N$ is an ordered list of points of \mathbb{E}^3 , and for any such p , there are $N(N-1)/2$ pairwise distances between the points in \mathbb{E}^3 that make it up. The symmetries of three-dimensional Euclidean space, \mathbb{E}^3 , are transformations—that is, one-one mappings of the

⁵There is some variance in the literature in how the term “configuration space” is used. Sometimes it is used for the space of ordered $3N$ -tuples of real numbers. We will retain the capitalization as a reminder that the term is being used as defined here, a usage that may differ from its usage in some other works.

space to itself—that preserve all distances between points. The symmetries of Configuration Space are transformations that preserve all of the $N(N-1)/2$ pairwise distances just mentioned. It should be clear that, given any symmetry of \mathbb{E}^3 , there is a corresponding symmetry of $(\mathbb{E}^3)^N$. That is, all the symmetries of 3-dimensional Euclidean space are built into the structure of $3N$ -dimensional Configuration Space.

On the other hand, it makes no sense to ask, of a pair of points p, q in the Configuration Space $(\mathbb{E}^3)^N$, representing two different configurations, what the distance is between those points in Configuration Space. If we pass from one configuration of planets to another, we can ask, for each of the N planets, how much its distance has changed, but on Configuration Space there is no privileged way to combine these to define distance between configurations.⁶

Things are a bit different if we consider the set of configurations of N classical objects when some or all of the objects lack any intrinsic characteristics that distinguish them from each other. Swapping of any two such objects does *not* yield a distinct physical state of affairs. The configuration space for N indistinguishable objects consists of *unordered* N -tuples of points in 3-dimensional Euclidean space. Its structure is a bit more complicated, as its topology is not the same as $(\mathbb{E}^3)^N$. See Goldstein et al. (2005), Maudlin (2013), and Chen (2017) for discussions of this point.

The $3N$ -dimensional Configuration Space $(\mathbb{E}^3)^N$ must be distinguished from $3N$ -dimensional Euclidean space, $\mathbb{E}^{(3N)}$. Both of these are $3N$ -dimensional spaces. But $\mathbb{E}^{(3N)}$ is like \mathbb{E}^3 , and unlike $(\mathbb{E}^3)^N$, in that, given a pair of points p, q in $\mathbb{E}^{(3N)}$ and a chosen unit of distance, it *does* make sense to ask what the distance between p and q is. Given a point chosen as an origin, $3N$ mutually perpendicular directions in $\mathbb{E}^{(3N)}$, and a line segment chosen as unit of length, we can coordinatize $\mathbb{E}^{(3N)}$ via an ordered $3N$ -tuple of real numbers. Thus, both $\mathbb{E}^{(3N)}$ and $(\mathbb{E}^3)^N$ have coordinatizations such that specification of $3N$ real numbers is required to specify a point. But the coordinatizations have different structure; one is an ordered $3N$ -tuple of real numbers, the other, an ordered N -tuple of ordered triples of real numbers.

⁶One could, of course, define metrics on Configuration Space in many ways, if such were wanted. Here's how to define a family of such metrics. Given two configurations of the solar system, we could consider, for each planet, the distance that the planet is displaced when passing from one configuration to the other. We could then define a distance, δ , between configurations, by taking the square of δ to be a weighted sum of the squares of the planetary displacements. But, in the absence of considerations that would imbue a metric of this sort with physical significance, there's no point in doing so.

The two spaces, $(\mathbb{E}^3)^N$ and $\mathbb{E}^{(3N)}$ have a lot of built-in structure. Some of it is metric structure, that is, structure having to do with assigning distances and areas and volumes. Some of it is topological structure, that is, structure that, unlike the metric structure, is preserved by all continuous transformations of the space. And some of it has to do with the distinction between curves in the space that are smooth and those that are not, between surfaces that are smooth and those that are not, between mappings of the space to itself that are smooth and those that are not, and between functions from the space to \mathbb{R}^N that are smooth and those that are not. This is the structure that is preserved by all smooth transformations of the space, and is called *differential structure*. A space that has topological structure and differential structure is called a *differentiable manifold*. It is n -dimensional if there is a smooth mapping between the space and \mathbb{R}^n (or else it can be covered by a family of local smooth mappings).⁷ $(\mathbb{E}^3)^N$ and $\mathbb{E}^{(3N)}$ are both $3N$ -dimensional differentiable manifolds, each of which has, in addition to differentiable and topological structure, some metric structure. We can also consider a bare $3N$ -dimensional differentiable manifold, with no built-in metric structure.

We now have three sorts of $3N$ -dimensional spaces to play with, which have very different intrinsic structure: Configuration Space $(\mathbb{E}^3)^N$, $3N$ -dimensional Euclidean space $\mathbb{E}^{(3N)}$, and a bare $3N$ -dimensional differentiable manifold. Which, if any, of these spaces is the space on which the wave functions of WFR, for a system consisting of N distinguishable particles, are defined?

Standard quantum-mechanical wave functions for a system consisting of N distinguishable particles are defined on the Configuration Space $(\mathbb{E}^3)^N$, and take advantage of some of the metric structure of that space for their definition (see secto 4, below). One would expect an Interpretive version of WFR to follow suit. This would be a Mild flavour, as this would mean that the $3N$ -dimensional fundamental space has built into it all the symmetries of three-dimensional Euclidean space. Some discussions of WFR (see, in particular, Albert 2015, 2019) suggest that the space on which its wave functions are defined has no built-in metric structure; all the metric structure there is, is the emergent, low-dimensional structure. On such a view, the fundamental space of WFR is just a $3N$ -dimensional differentiable manifold. This would

⁷A differentiable manifold can be such that its dimension varies from place to place in the manifold, but we will not have to deal with manifolds like that, so mention of this point may be safely relegated to a footnote.

be a *Very Spicy* flavour of WFR. But it could not be merely Interpretive, as quantum-mechanical wave functions lean heavily on metric structure of the spaces on which they are defined for their very definition.

3 Preliminaries: Classical fields on space-time

Proponents of WFR often draw an analogy between classical fields and wave functions. Quantum wave functions are like classical fields in some respects, and unlike them in others. One area of similarity has to do with the dependence on background spacetime structure of a mathematical representation of a classical field or a quantum wave function. A specification of a classical field that permits one to talk about the intensity of the field at a point requires one to invoke features of the metric structure of the space on which it is defined; a mere differentiable manifold of the appropriate dimension does not suffice. Quantum wave functions are no different in this respect. To make this clear, it is helpful to bear in mind the way in which the assignment of classical field values to points in space relies on metric structure of the background space. None of what follows will be controversial, or, to most of my readers, novel. We are, as Wittgenstein would put it, “assembling reminders for a particular purpose” (Wittgenstein, 1953, §127).

Think, first, of a mass-density or charge-density field on 3-dimensional Euclidean space. It is the job of such a field to represent a mass or charge *distribution* in space. A mass distribution associates with any region of space you can name a quantity of mass in that region. Obviously, we can’t specify a quantity of mass just by giving a number; we have to say also what units we’re using—grams, kilograms, or slugs, for example. Given a choice of unit of mass—say, kilograms—we have numbers assigned to regions of space, representing the mass, in kilograms, contained in those regions. What remains the same, for any choice of units, is the *ratio* of the numbers assigned to any pair of regions. That’s where the real physical content lies; a choice of unit is a choice of a standard mass to which we assign the number one. When we specify the mass of an object by specifying a number relative to some unit, that number is the ratio of the mass to the mass of the chosen unit. All this is obvious, but it will serve as well to bear in mind in what follows: whenever we use some bit of math to represent something physical, the physical content of our representation is the stuff that doesn’t vary according to our arbitrary choices.

Provided that, at any point p in space, the mass contained in a region that includes p shrinks to zero as the volume of the region shrinks to zero, we can represent a mass distribution by a *mass density function* ρ .⁸ The way that the mass density function represents a mass distribution is: for any region R that is assigned a mass by the mass distribution, the mass contained in R is the integral of the mass density over that region.

The number that a mass density function assigns to a point in space depends on two things: choice of unit of mass, and choice of unit of volume. The same mass density might be 1 gram per cubic centimetre, or 1,000 kilograms per cubic metre. If we have two different mass distributions, represented by two mass density functions ρ_1 and ρ_2 , the *ratio* of ρ_1 to ρ_2 at a given point will be the same (at least, for almost all points; exceptions on a set of total volume zero don't make a difference to the mass distribution represented) no matter what unit of volume we're using, and so it's those ratios that have physical significance.

You might have been tempted to say that the ratio of $\rho_1(p)$ to $\rho_1(q)$, for two different points, has physical significance on its own, but that's not right, without further qualification.

Consider the following mass distribution. In each of two nonoverlapping cubical regions of space, R_1 and R_2 , there is half a kilogram of mass, and none anywhere else. Within each of these regions the mass is uniformly distributed, meaning that the amount of mass in any subregion of R_1 or R_2 is proportional to the volume of the subregion. The sides of the region R_2 are twice as long as the sides of R_1 , and hence the volume of R_2 is eight times the volume of R_1 .

Let $V(R_1)$ and $V(R_2)$ be the volumes of the two regions of nonzero mass. We can define a density function ρ_A for this mass distribution that is equal to $1/(2V(R_1))$ in the interior of R_1 , to $1/(2V(R_2))$ in the interior of R_2 , and is zero everywhere else. Since the same amount of mass is found in both R_1 and R_2 , and R_2 is larger than R_1 , the mass is more thinly spread out in R_2 than it is in R_1 , and our mass density function reflects this; it takes on smaller values in R_2 than it does in R_1 .

⁸Distributions that do not have this property, which are said to be *singular* with respect to the volume measure, can be represented by elements of a class of mathematical objects called *generalized functions*, which generalize the notion of a density function. For the sake of illustration, we consider a non-singular distribution, which has a density function. All the conceptual points we'll make remain intact if one expands the class of distributions considered to include singular distributions.

All of this is taking place within Euclidean space, and our characterization of the mass density function leans heavily on the structure of that space. The two regions R_1 and R_2 contain the same total mass, evenly distributed with the regions, and R_2 has eight times the volume of R_1 , and hence the mass density within R_1 is larger than it is in R_2 . To say all of this, we are relying on the fact that we are working within a space that comes with a notion of rigid displacement, and hence a standard way of comparing volumes of distinct regions of space.

But we might also consider some nonstandard scheme of assigning volumes to regions. Suppose that someone (call him Bob), uses, instead of standard rulers, rulers that change their size when moved from place to place. Suppose that, when transported from R_1 to R_2 , the rulers quadruple in size. For simplicity, suppose that they don't change their size when moved about within these two regions. Judged by Bob's rulers, the sides of R_2 are *half* the length of the size of those of R_1 , and therefore, on this way of measuring, the volume of R_2 is one-eighth of the volume of R_1 . Bob will judge the mass to be more spread out in R_1 than it is in R_2 , and hence will write down a mass density function, ρ_B , that takes on a smaller value at points in R_1 than it does at points in R_2 .

So, we have two mass density functions, one defined with respect to a standard way of assigning volumes to regions of space, and another defined with respect to Bob's eccentric system of mensuration. These two mass density functions *agree on the amount of mass present in any region of space*. This illustrates the fact that a mass density function represents a mass distribution only in conjunction with some way of measuring volume. When Bob uses his mass density function to calculate the amount of mass in a given region, he will of course use his own measure of volume to do the integration. Our original mass density function, paired with the standard way of measuring volumes, and Bob's mass density function, paired with Bob's eccentric system of assigning volumes to regions of space, represent the same distribution of mass.

This is not a case of underdetermination of theory; we have two modes of expression, expressing the same physical facts. Recall that it's the job of a mass density function to represent a mass distribution. The two density functions agree on the amount of mass assigned to any region of space; that is, they agree on what they're saying about the distribution of mass in the world. This illustrates a general point. If you want to know when a difference in the mathematics used to represent some physical state of affairs corresponds to

a difference in the physical state of affairs represented, and when it doesn't, simply staring at the mathematics will be of no use. There is no purely formal criterion. You have to think about what the mathematics is being used *for*.⁹

How different can two mass density functions representing a given mass distribution be, subject only to the constraint that they represent the distribution with respect to *some* measure on Euclidean space, that is, *some* systematic way of assigning volumes to regions of space? A little reflection shows that there is very little restriction. Suppose we restrict ourselves to considering only measures that assign zero volume to every region that has zero volume on a standard, Euclidean measure. Then, if one density function is zero everywhere within some region of space, another density function must be zero almost everywhere within that region (that is, nonzero within that region on at most a set of points of total volume zero). But this is the only constraint; subject to this restriction, anything goes, and *any* function that is integrable with respect to Euclidean measure can represent *any* given mass distribution.

One might be inclined to say that the mass density function that we originally gave, defined with respect to the usual notion of volume on Euclidean space, is a respectable one, and that Bob's is somewhat disreputable. There's something right about this, of course, in that the volume measure used for the original density function respects the symmetries of Euclidean space, assigning the same volume to sets related by rigid transport. Another way of saying this is that it is what is measured by measuring rods that we usually take (perhaps on the basis of considerations concerning the physics of such rods) to not change length when transported from place to place. Our ability to say this depends, of course, on the fact that we are operating within a space that has symmetries worth respecting! But, even if one grants that one of these mass density functions is respectable and the other disreputable, it bears repeating there is *no sense* in which one of these density functions is *correct* and the other *incorrect*, as they *don't disagree* on what they say about the distribution of mass in space.

Exactly the same considerations apply, of course, to other sorts of density functions, including charge density functions representing charge distributions, and probability density functions representing probability distributions. These are defined with respect to some way of assigning volumes to

⁹File this under *Things that should go without saying but unfortunately don't*.

regions of the space on which they take their values—that is, to what mathematicians call a *measure*. And this means that similar considerations apply to electromagnetic fields. Maxwell’s equations, in differential form, relate electric and magnetic fields to charge and current densities. Since the charge and current densities are defined relative to a measure on the background space, and change upon change of measure used, the same must be true of electric and magnetic fields. Another way to see this is that the energy density associated with electromagnetic radiation is proportional to the square of the field amplitude.

There’s another way in which electric and magnetic fields are dependent on background spacetime structure. The force on a charged particle exerted by electromagnetic fields can be partitioned into a component proportional to the electric field and independent of the velocity of the particle, and a component that involves a product of the magnetic field and the particle’s velocity. The electric field is, therefore, proportional to the force on a charged particle at rest, that is, a particle with velocity equal to zero. Talk of velocity of a particle only makes sense once a reference frame is chosen, to be used as a standard of rest. This means (as Einstein was perhaps the first to really see clearly) that a partition of the electromagnetic field into an electric field and a magnetic field is relative to a reference frame. It makes no sense to ask what *the* value of an electric field is at some place and time, but only what its value is relative to some reference frame.

Most of this will be familiar to many readers. Why am I rehearsing it? Because it is commonly said, by proponents of WFR, that quantum wave functions are like classical fields, and also that they assign *numbers* to points in the fundamental space. But a classical field doesn’t yield a number until we specify some units of measurement (and for that, we need to know what sort of physical quantity the field is—a mass density? a charge density? an electric field?) and also a measure on the space on which it is defined. As we will see, quantum-mechanical wave functions are like classical fields in this respect—they are defined with respect to a measure on the underlying space.

4 The plurality of wave function representations of a quantum state

The considerations of the previous section, that we brought to bear on classical fields, apply with equal force to quantum-mechanical wave functions defined on Configuration Space. This becomes obvious when we consider that, for any wave function ψ , its squared absolute value, $|\psi|^2$, is a density, standardly interpreted as a probability density function.

As this fact—not an abstruse one, but built into the structure of elementary quantum mechanics—about the relativity of a quantum wave function to a background measure on the space on which it is defined is not always emphasized in the philosophical literature, it is worthwhile to go through considerations of the sort discussed in the previous section, and apply them to wave functions.

For simplicity, let's discuss the case of a pair of distinguishable spinless particles. Let R_1 and R_2 be regions of space as above, and let S_1 and S_2 be two nonoverlapping cubical regions of space that are congruent to each other. We imagine an experiment involving particle “detectors” capable of distinguishing the two particles, arrayed over all of space.¹⁰ Consider a quantum state that yields probability $1/2$ for “detection” of particle 1 in R_1 and particle 2 in S_1 , and probability $1/2$ for “detection” of particle 1 in R_2 and particle 2 in S_2 , with uniform probability within each region. Let Ω_1 be the subset of the Configuration Space of a pair of particles, consisting of pairs of points (p_1, p_2) with p_1 in R_1 and p_2 in S_1 , and let Ω_2 be the subset of Configuration Space consisting of pairs of points with p_1 in R_2 and p_2 in S_2 . Let ψ_A be a wave function that is zero everywhere except in Ω_1 and Ω_2 , with constant value $1/\sqrt{2V(R_1)V(S_1)}$ in Ω_1 , and constant value $1/\sqrt{2V(R_2)V(S_2)}$ in Ω_2 . Because the probabilities for Ω_1 and Ω_2 are equal, and Ω_2 is larger than Ω_1 , the wave function is more “spread out” in Ω_2 —that is, the amplitude of ψ_A is smaller in Ω_2 than it is in Ω_1 .

Bob can also represent this quantum state, using his strange length-changing rulers, to measure distances (and hence volumes of space). Recall that, on Bob's way of measuring things, the region R_2 has a volume that

¹⁰The scare-quotes around “detectors” are there to flag the fact that we don't want to presume that firing of a “detector” indicates a pre-existing location of a particle. Take talk of “detection” to mean: an appropriate piece of experimental apparatus, of the sort usually called a particle detector, yields a positive result.

is one-eighth that of R_1 . Bob writes down a wave function ψ_B that has a constant value $1/\sqrt{2V_B(R_1)V_B(S_1)}$ in Ω_1 , and constant value $1/\sqrt{2V_B(R_2)V_B(S_2)}$ in Ω_2 , where V_B is volume measured by Bob's rulers. Suppose that Bob's rulers don't change length when moved from S_1 to S_2 . Then, since, by Bob's measurements, the volume of R_2 is smaller than that of R_1 , his wave function is larger in Ω_2 than it is in Ω_1 .

The two wave functions ψ_A and ψ_B represent the same quantum state. This is a point that deserves emphasis. *Though they differ as to whether the amplitude is larger in Ω_1 or in Ω_2 , the two wave functions, ψ_A and ψ_B , each in its own way, represent the same quantum state.* There is no such thing as the wave function that represents a given quantum state. A Configuration Space wave function does not represent a quantum state on its own, but only in concert with a measure on Configuration Space.

As with the mass densities, one of these wave functions may be regarded as more respectable than the other, on the grounds that one of them, and not the other, is defined with respect to a measure that respects the symmetries of the space on which it takes its values—in this case, the Configuration Space of a pair of distinguishable particles, which is the set of ordered pairs of points in three-dimensional Euclidean space, and which has all the symmetries of three-dimensional Euclidean space.

If the underlying space is Configuration Space—that is, the set of all ordered N -tuples of points in Euclidean Space—we can point out that the measure that Bob has used to define his wave function employs a measure that doesn't respect the symmetries of our underlying space, and so regard Bob's wave function as somewhat disreputable. If, however, the fundamental space of WFR has no built-in relation to Euclidean 3D space, and no other metric structure that yields a notion of rigid displacement, this move is unavailable.

If I have a wave function on Configuration Space defined with respect to some measure, that I am using to represent some quantum state, what restrictions are there on wave functions, that I might choose to define with respect to some other measure, to represent the same quantum state? The same considerations that we applied to mass densities in the previous section. Suppose I have a wave function ψ on Configuration Space that I am using, in concert with some measure on Configuration Space, to represent some quantum state, and I want to change the amplitude of the wave function. There is little restriction on this; any function—subject only to the restriction that it be equal to zero at the same points that ψ is—can be the amplitude of

a wave function that, in concert with a suitable choice of measure, represents the same quantum state that ψ does. As with the mass density functions, there is no question of one being correct and the other incorrect. The job of a wave function is to represent a quantum state, and these two functions *represent the same quantum state*.

Something similar can be said about the phase of the wave function. There is another source of freedom in choosing a wave function to represent a given quantum state, which is called *gauge freedom*. It is usually not mentioned in quantum mechanics textbooks until the coupling of charged particles to applied magnetic fields is discussed, but it is there nonetheless. We leave discussion of this to the next section, where we outline the full range of wave function representations of quantum states.

Flexibility in choice of measure used, together with gauge freedom, means that, given any quantum state, if we want to represent that quantum state via a wave function on Configuration Space, we can use any function we want, provided that it is equal to zero at all (or almost all) points in any region of Configuration Space to which the quantum state assigns probability zero.¹¹

There is also a relativity of wave functions to reference frames, somewhat reminiscent of the frame-dependence of electric and magnetic fields. Let's restrict ourselves to ordinary, non-relativistic quantum mechanics, in Galilean spacetime. Given a scheme for representing quantum states by wave functions on Configuration Space, one and the same quantum state will be represented by different functions, with respect to different reference frames. The effect of transformations from one reference frame to another leaves the amplitude at any given point the same, but changes the phase in a way that varies from point to point (see Levy-Leblond 1963; Brown and Holland 1999).

5 The cornucopia of wave function representations of quantum states

In the previous section we saw that, even if one privileges wave functions on Configuration Space, there is a multitude of wave function representations of any quantum state. But Configuration Space wave functions are not the only wave functions. In this section we discuss the full range of wave function representations of a given quantum theory and its states. A wave function

¹¹See Appendix for a more detailed exposition of this point.

representation of a quantum theory is one kind of Hilbert space representation; the functions (or, rather, equivalence classes of them) are themselves vectors in a Hilbert space. This is perhaps worth emphasizing, because in the WFR literature one sometimes gets the impression that Hilbert space representations are being *contrasted* with wave function representations. We will assume that the reader knows what a Hilbert space is. If you don't, pause and consult Ismael (2015). We'll wait.

In quantum mechanics, the notion of a *complete set* of dynamical variables (usually called a complete set of *observables*) plays a key role. Roughly speaking, a set of variables is a complete set if a specification of definite values for all of them uniquely determines the quantum state of the system. For example, for a system of N distinguishable particles without spin, positions of all the particles form a complete set, as do momenta of all the particles, or positions of some and momenta of others. If they have spin, then specification of positions of all particles does not suffice, as it leaves open their spin values.

The notion of *compatible* observables also plays a key role. Two observables are compatible if quantum theory places no limitations on how precisely they may be simultaneously defined by a quantum state (unlike, say, position and momentum of a particle, or components of spin in different directions).

To construct a wave function representation of a quantum theory, one selects a complete set of compatible variables $\mathcal{A} = \{A_1, A_2, \dots, A_n\}$. The set of possible values a variable can take on is known as its *spectrum*. A joint specification of values of all of the variables in \mathcal{A} consists of selecting, for each of those variables, a point in its spectrum, and can be represented by an ordered n -tuple (a_1, a_2, \dots, a_n) . The set of all possible n -tuples of this sort is the *joint spectrum* of the variables in \mathcal{A} . Call this set $\Sigma_{\mathcal{A}}$. It is the space on which our wave functions will take values.

We want to construct a Hilbert space out of functions on $\Sigma_{\mathcal{A}}$.¹² So, the next step is to choose a measure μ on $\Sigma_{\mathcal{A}}$, that is, a way of assigning

¹²There's a wrinkle, which is of minor importance, and so can be relegated to a footnote. Though this is not always mentioned in introductory quantum-mechanical textbooks, there is a difference between functions on $\Sigma_{\mathcal{A}}$ and vectors in a Hilbert space. For any two vectors u, v in a Hilbert space, if the norm of their difference, $\|u - v\|$, is zero, then $u = v$. If we define a norm in terms of integration with respect to some measure, this norm, applied to two functions ψ, ϕ , might be zero, even if the two functions are not identical, because they might differ only on a set of total measure zero. For this reason it is usually said that the elements of our Hilbert space are not functions, but *equivalence-classes* of functions that differ at most on a set of measure zero.

a volume, in a generalized sense, to subsets of $\Sigma_{\mathcal{A}}$. This will be used to define an inner product on the set of square integrable functions on $\Sigma_{\mathcal{A}}$. We have some freedom in choosing this measure, provided only that it not assign measure zero to any subset of $\Sigma_{\mathcal{A}}$ that is assigned nonzero probability by some quantum state. This flexibility in choice of measure is not a merely otiose mathematical curiosity. In proofs that function representations of Hilbert spaces exist in the first place, one typically chooses a measure defined by an appropriate vector to construct one representation. Then, if one likes, one can make a transformation to a representation in terms of some other measure (see Reed and Simon 1980, §VII.2 for discussion). In addition, sometimes what one is interested in is how a given state differs from another state, say, the ground state of the system. One might then define wave functions with respect to a measure induced by that state.

This fixes a correspondence between the dynamical variables in the set \mathcal{A} and operators on the Hilbert space of functions on $\Sigma_{\mathcal{A}}$. Corresponding to a variable A_i is an operator \hat{A}_i that is just a multiplication operator. That is, the effect of the operator \hat{A}_i , corresponding to physical variable A_i , on a function $\psi(a_1, a_2, \dots, a_n)$ is to multiply the function by a_i . Similarly, for any function $F(A_1, A_2, \dots, A_n)$ of these variables, the corresponding operator on the space of functions is the one that multiplies a function by $F(a_1, a_2, \dots, a_n)$.

This doesn't yet fix a wave-function representation of our quantum theory. In the Hamiltonian formulation of classical mechanics, dynamical variables come in pairs, consisting of a coordinate and its conjugate momentum. For example, if the coordinate is an ordinary space coordinate, say, the x -coordinate (relative to some coordinate system) of the position of a particle, one can take, as momentum conjugate to that coordinate, ordinary linear momentum in the x -direction, p_x . If the coordinate is an angle, one can take angular momentum as conjugate momentum. What can be taken as a coordinate is fairly general (for this reason, one sees talk of *generalized coordinates*), and the distinction between coordinates and their conjugate momenta is not absolute. We could, for example, take as generalized coordinates the components, with respect to some system of spatial coordinates, of some body's momentum, $\{p_x, p_y, p_z\}$, and $\{-x, -y, -z\}$ as momenta conjugate to these generalized coordinates.

Quantum theory imposes relations, called the *canonical commutation relations*, on operators corresponding to a coordinate and its conjugate momentum. These are as follows. Given generalized coordinate-momentum pairs

$\{(q_i, p_i)\}$, to be represented by operators $\{(\hat{q}_i, \hat{p}_i)\}$, we require:

- The operators $\{\hat{q}_i\}$ all commute with each other, as do all the operators $\{\hat{p}_i\}$. What this means is that, for any i, j ,

$$\begin{aligned}\hat{q}_i \hat{q}_j &= \hat{q}_j \hat{q}_i; \\ \hat{p}_i \hat{p}_j &= \hat{p}_j \hat{p}_i.\end{aligned}$$

- For distinct i, j , \hat{q}_i commutes with \hat{p}_j :

$$\hat{q}_i \hat{p}_j = \hat{p}_j \hat{q}_i.$$

- The operators corresponding to a coordinate and its conjugate momentum satisfy,

$$\hat{q}_i \hat{p}_i - \hat{p}_i \hat{q}_i = i\hbar \hat{I},$$

where \hat{I} is the identity operator, and \hbar is Planck's constant, h , divided by 2π .

Before we have a function-space representation of our quantum theory, we have to choose operators on the space of functions to represent the momenta conjugate to our selected variables $\{A_1, A_2, \dots, A_n\}$, satisfying the canonical commutation relations.

Though this typically doesn't get mentioned in the introductory sections of a quantum mechanics textbook, we have some freedom in doing so. For concreteness, suppose that the configuration of a system is represented by coordinates $(x_1, x_2, \dots, x_{3N})$ that can each take the full range of real numbers as values. We want to construct a configuration-space wave-function representation of our quantum theory, that is, a representation on which any pure quantum state can be represented by a function $\psi(x_1, x_2, \dots, x_{3N})$. The operators $(\hat{x}_1, \hat{x}_2, \dots, \hat{x}_{3N})$ corresponding to the position degrees of freedom are multiplication operators:

$$\hat{x}_i \psi(x_1, x_2, \dots, x_{3N}) = x_i \psi(x_1, x_2, \dots, x_{3N}). \quad (1)$$

It is easy to check that, for any smooth function $\alpha(x_1, x_2, \dots, x_{3N})$, the operators $(\hat{p}_1, \hat{p}_1, \dots, \hat{p}_{3N})$ defined by

$$\hat{p}_i \psi = -i\hbar \frac{\partial \psi}{\partial x_i} + \left(\frac{\partial \alpha}{\partial x_i} \right) \psi \quad (2)$$

satisfy the canonical commutation relations. Moreover, different choices of the function α yield unitarily equivalent representations of our quantum theory.¹³ In the literature on such things, there is universal agreement that unitarily equivalent representations of an algebra of observables are physically equivalent. The transformation from a representation with one choice of α to a representation with another choice changes the phase of the wave function, in such a way that, for any quantum state, we can have the phase of the corresponding wave-function be any smooth function on Configuration Space that we want it to be.

All of the wave-function representations, for different choices of α , and hence different choices of operators to represent the momenta $\{p_i\}$, are perfectly cromulent wave-function representations of one and the same quantum theory. Here again, it should be emphasized: this is not a case of underdetermination of theory; it's a matter of choice of mathematical machinery to represent quantum states of one and the same theory.

So, to sum up: to get a wave function representation in which all the quantum states of some system at some time t_0 are represented by wave functions, we need:

1. A choice of complete set of variables (“observables”), $\mathcal{A} = \{A_1, A_2, \dots, A_n\}$. The set $\Sigma_{\mathcal{A}}$ of n -tuples (a_1, a_2, \dots, a_n) of points in the spectra of these observables will be the space on which our wave functions are defined.
2. A measure μ on $\Sigma_{\mathcal{A}}$, which allows us to construct a Hilbert space whose elements are equivalence classes of functions that are square-integrable with respect to μ , using μ to define the inner product for this Hilbert space. The measure can be any measure that does not assign measure zero to any region of $\Sigma_{\mathcal{A}}$ that is assigned nonzero probability by some quantum state.
3. A choice of operators on this space, satisfying standard commutation relations, to represent observables that are not functions of the variables in the chosen set \mathcal{A} . This will include the variables conjugate to the chosen variables.

Even if we restrict ourselves to wave functions on Configuration Space, for any quantum state Ψ :

¹³See Appendix for definition, if this is not familiar, and for discussion of this claim.

1. The freedom we have to choose the measure used to define the inner product for our wave-function representation means that we can choose the amplitude of the wave function ψ used to represent Ψ to be any function we want, with one caveat: If there is a subset Δ of Configuration space that is assigned probability zero by Ψ and is assigned nonzero probability by some other state, ψ should be zero almost everywhere on Δ .
2. The freedom we have to choose the operators representing the momenta means that we have complete freedom to choose the function that is the phase of the wave function ψ that we use to represent Ψ .

Changing the wave function used to represent a given quantum state changes also the wave function used to represent any other quantum state. Once we choose the dynamical variables that our wave functions are to be functions of (choosing, for example, a Configuration Space wave function representation), there is something that all representations have to have in common. For any two quantum states Ψ_1 and Ψ_2 , different wave-function representations will agree on the *ratio* of the corresponding wave functions, that is, on the ratio of $\psi_1(\mathbf{x})$ to $\psi_2(\mathbf{x})$, for all (or at least almost all) points \mathbf{x} where $\psi_2(\mathbf{x})$ is nonzero. This, as you may recall, is similar to the conclusion we came to in connection with mass densities.

6 Dynamics

The previous section had to do with representations of the state of a system at a given time. To represent the state of a system as it evolves with time, we need Hilbert space representations for the state at different times. This requires more choices to be made, which gives rise to a further plurality of representations.

Suppose we have some system evolving in time. Suppose, further, that we regard some experiments performed at different times as, in an important sense, the same experiment, measuring the same observable, at different times. This could be, for example, a measurement of position with respect to some reference frame. (Note that this depends on choice of reference frame; what counts as “same position” at different times will vary with reference frame.)

Suppose that we have a Hilbert space representation of the state of the system at some time t_0 . We want representations at other times, too. There are two main choices:

- *Schrödinger picture.* We use the same operator for what we are counting as the “same observable” at different times. As the state changes, this means using different vectors to represent the state at different times.
- *Heisenberg picture.* We use the same state vector to represent the state at all times, and use different operators to represent the “same observable” at different times.

There are, of course, other possible choices, but these are the two most commonly used in nonrelativistic quantum mechanics.

These are two different ways to construct Hilbert space representations of evolving states of a single quantum theory. As they agree on probabilities of outcomes of all possible experiments, they count, in standard quantum mechanics, as physically equivalent.

One sometimes hears it said that, on the Heisenberg picture, the state is unchanging. This is false. An unchanging state would mean that the probabilities of outcomes of an experiment are independent of the time at which the experiment is performed. What is unchanged is the vector used to represent the state. But we should always remember that a Hilbert state vector represents a quantum state only in concert with some choice of ways to associate operators on the Hilbert space with physical observables.

7 The fundamental arena as a differentiable manifold

A recurring theme in David Albert’s writings on WFR is that three-dimensional geometrical structure is to be thought of emergent from the behaviour of the wave function. That is, though there is no three-dimensional geometric structure built into the structure of the fundamental space or into the fundamental laws governing the behaviour of the wave function, it could be that there is a coordinatization $(x_1, y_1, z_1, \dots, x_N, y_N, z_N)$ of the fundamental space such that the behaviour of the wave function can be simply and

informatively encapsulated by a law of evolution that depends only on the $N(N-1)/2$ pairwise 3D distances,

$$d_{ij} = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2 + (z_i - z_j)^2}. \quad (3)$$

Here's another way to put the point. Let Γ be the fundamental space. The hypothesis is that there are $3N$ smooth functions from Γ to the real numbers, which we will call $(x_1, y_1, z_1, \dots, x_N, y_N, z_N)$, which are such that no two points in Γ have the same values of all these functions, and which are such that the evolution of the wave function may be simply described in terms of the functions d_{ij} defined in terms of them.

For instance, the history of the wave function might be simply and informatively encapsulated by saying that it enacts a GRW evolution with respect to these coordinates. That is, the wave function is, at random moments, multiplied by a function of the form

$$f_i(x_1, y_1, z_1, \dots, x_N, y_N, z_N; a_i, b_i, c_i) = e^{-\alpha((x_i - a_i)^2 + (y_i - b_i)^2 + (z_i - c_i)^2)} \quad (4)$$

To define the probability rule for the “locations” (a_i, b_i, c_i) of these hits, we first use the coordinatization $(x_1, y_1, z_1, \dots, x_N, y_N, z_N)$ to define a measure on the fundamental space, via

$$d\mu = dx_1 dy_1 dz_1 \dots dx_N dy_N dz_N. \quad (5)$$

Given a wave function ψ , we then define a probability density for the parameters (a_i, b_i, c_i) ,

$$p(a_i, b_i, c_i) \propto \int_{\Gamma} |f_i(\mathbf{x}; a_i, b_i, c_i) \psi(\mathbf{x})|^2 d\mu. \quad (6)$$

This yields probabilities for the location of the hits, when the wave function is ψ .

In between hits, the wave function evolves according to a Schrödinger equation of the form,

$$i\hbar \frac{\partial \psi}{\partial t} = - \sum_i \frac{\hbar^2}{2m_i} \left(\frac{\partial^2}{\partial x_i^2} + \frac{\partial^2}{\partial y_i^2} + \frac{\partial^2}{\partial z_i^2} \right) \psi + \sum_{i,j} V_{ij}((x_i - x_j)^2 + (y_i - y_j)^2 + (z_i - z_j)^2). \quad (7)$$

A dynamical law of this form is invariant under the set of transformations that leave all of the functions d_{ij} invariant—that is, a group of transformations isomorphic to the rigid transformations of 3D Euclidean space.

Albert calls the distances d_{ij} *interaction distances*. The space of possible interaction distances has the structure of 3D Euclidean space. As it is the interaction distances that play a role in the dynamics, it is these distances that enter into our descriptions of material objects.

what it is to be a table or a chair or a building or a person is—at the end of the day—to *occupy a certain location in the causal map of the world*. The thing to keep in mind is that the production of geometrical appearance is—at the end of the day—a matter of *dynamics*. (Albert, 2015, 127)

He concludes (rightly, in my opinion) that any intrinsic geometrical structure that the background space might have plays no physically significant role.

the affine and metrical structure of the background, fundamental, pre-dynamical space does *no explanatory work whatsoever* in this story—and that once that structure is *dismantled*, the very *idea* of anything like a ‘Fundamental Space of the World’ disappears along with it. (2019, 93)

Albert distinguishes between two structures that, in pre-quantum physics, had been thought to coincide: a *fundamental arena* of the world, and the physically significant geometrical space.

Think of the fundamental arena as a set of points that amounts to something like the *totality of opportunities for things, at any particular time, to be one way or another*. Or you could put it this way: what we have in mind, what we mean to say, when we refer to some set of points as the fundamental arena of the world, is that a specification of what is physically going on, at each one of those points, at any particular time, amounts to a *complete specification of the physical situation of the world* at that time. (2019, 89–90)

The fundamental arena on which the history of the world unfolds is a high-dimensional manifold, but it has no intrinsic metric or affine structure.

The only conception of *distance* that does any dynamical or predictive or explanatory *work* in the theory, . . . , is the 3-dimensional Pythagorean distance $(x_k - x_j)^2 + (y_k - y_j)^2 + (z_k - z_j)^2$. And it deserves to be emphasized that there is absolutely nothing approximate or defective or misleading or illusory or otherwise second-class—on a picture like this one—about our everyday experience of the geometry of the world as Euclidian and 3-dimensional. On the sort of picture I have been sketching here, the Euclidian 3-dimensional geometry of our everyday experience—notwithstanding that it is something *emergent*—is (again, and indeed, and on the contrary) the *true* and *unique* and *authentic* and *exact* and *complete* geometry of the world. Period. End of story. (94)

I am in wholehearted agreement with Albert’s remarks on geometric structure, and I myself have argued that structure that plays no dynamical role just isn’t metric structure (Myrvold, 2019).

So, the fundamental arena is to be thought of as a bare differentiable manifold, with no intrinsic metrical or affine structure.

But now that we are dealing with a quantum-mechanical, field-like *wave-function*, the thought is that a fundamental differential manifold, with no affine or geometrical structure at all, will suffice. (Albert, 2019, 9, fn. 8)

This poses a dilemma, however. As we have seen, both classical fields and quantum-mechanical wave functions require something more than topological and differential structure for their definition: they require a measure on the space on which they are defined, which assigns volumes to regions of that space, and different choices of measure yield different functions. Either the wave functions of which Albert speaks are defined with respect to some measure, or else they are importantly unlike quantum-mechanical wave functions.

One way to go would be to postulate that the fundamental arena is a differentiable manifold with a *fundamental measure* defined on it.

But wait. Once we have the coordinatization of which Albert speaks, which yields interaction distances d_{ij} , which are not intrinsic structure of the fundamental arena but are picked out by the way that things have happened to unfold, we can use them to define the measure (5). This measure *will* have physical significance, because, for one thing, the integral of $|\psi|^2$ with

respect to this measure will be conserved as long as the wave function is evolving according to (7), and, for another, it is integrals with respect to this measure that are used to define probabilities for the GRW jumps. And everything that Albert says about background metric structure applies to the alleged fundamental measure. Any background measure used to define the wave function, if it is not the one defined by (5), will play no role in the dynamics, and will not be physically meaningful.

So, we have here a question about the nature of the function of which Albert speaks, and what it is that it assigns to points in the fundamental arena. If it's anything like a quantum-mechanical wave function, it is defined with respect to some measure, and what function it is will vary (pretty much arbitrarily!) with choice of measure.

A fundamental measure is not in the spirit of Albert's overall approach. The best way, I think, to construe Albert's proposal is to retain the idea that the fundamental arena has no intrinsic structure other than topological and differential structure, and to take it that the function that he calls the wave function is *not* anything like a classical field or a quantum-mechanical wave function. It is frequently said, in the WFR literature, that a wave function assigns a *complex number* (or, equivalently, two real numbers, an amplitude and a phase) to a point in the fundamental arena. I think we should take that literally. On this proposal, there is a matter of physical fact about what number the function of which Albert speaks assigns at any given time to any point in the fundamental area, and this really is just a *number*. Its square does not represent a probability density or a stuff-density or any kind of density at all, because the value of a density is defined only relative to some measure; it's not a something-or-other per unit volume, but just a number.

This pure-number valued field would be something *sui generis*, something very much unlike either quantum wave functions or classical fields. That isn't, in itself, an objection. It is not an objection to classical electrodynamics, in the modern conception of the subject, that electromagnetic fields are unlike anything in classical mechanics, and it is not an objection to quantum mechanics that quantum wave functions are unlike classical fields or anything else in pre-quantum physics. But it is true that, on this proposal, wave function realism should be thought of as a thoroughly Constructive project. We have come a long way from the original suggestion that it is simply the view that emerges from taking quantum mechanics realistically.

8 Wave function realisms

At least four questions must be addressed, before even one version of wave function realism has been formulated, even for the simplest case of the non-relativistic theory of a finite number of spinless particles.¹⁴

1. What is the structure of the fundamental space?
2. What is it that a wave function assigns to points in this space?
3. Do different wave functions always correspond to physically distinct states of affairs?
4. What dynamical laws are assumed, for wave functions?

One answer to the first question would be to take literally a phrase that occurs repeatedly in the literature, that the fundamental space has the structure of a classical Configuration Space. This would mean that we have all the structure needed to define standard quantum wave functions, which, as we have seen, rely heavily on background spacetime structure. It would also mean that the symmetry group of the fundamental space is the same as the symmetry group of the low-dimensional space; the structure of the low dimensional space is built into the structure of the high-dimensional fundamental space. This would be a flavour of WFR that is Very Mild Indeed.

A Spicy version of WFR might have it that, as David Albert suggests, the fundamental space is a bare $3N$ -dimensional manifold, with no structure other than its topological and differential structure. As we have seen, this is insufficient for defining a quantum-mechanical wave function, which requires, at minimum, a measure on the joint spectrum of the set of observables chosen to define the function representation of the state.

A phrase that often occurs in the literature on WFR is “the wave function,” suggesting that, for any way the world could be, there is such a thing as *the* wave function that represents the way that things are. In standard quantum mechanics there is no such thing. As others before me have pointed out, in standard quantum mechanics, there are multiple wave functions that can be used to represent a single quantum state. One need not privilege Configuration Space wave functions, and in quantum mechanics one routinely employs other representations, such as a momentum-space representation.

¹⁴The reader can probably think of more.

Even if we decide that Configuration-Space wave functions have a privileged status, for any quantum state there are multiple wave functions, as multiplying a wave function by a constant yields a new function that represents the same quantum state. This doesn't exhaust the range of Configuration Space wave functions. A Configuration-Space wave function represents a quantum state only in conjunction with some measure on Configuration Space, and some choice of operators to represent the conjugate momenta. *Any* function can represent a given quantum state, provided that only that it is zero in the right places.

It is not always clear whether talk of “the wave function” is to be taken literally. One option for the Constructivist wave function realist, willing to depart radically from standard quantum mechanics, is something along the following lines. There is, indeed, at the fundamental level a unique wave function that represents the physical state of the world, and it assigns a pure number to points in the fundamental arena, not defined with respect to some background measure. All the other wave functions in the cornucopia are at best calculational tools. A physicist is, of course, free to use whichever wave function suits a given computational task, but at most one of the many wave functions that, in standard quantum mechanics, can all represent the same quantum state, represents what reality is like at the fundamental level. This would, of course, be a heavily Constructivist position. The underlying theory, though it might recover the predictions of standard quantum mechanics, would have a structure very different from standard quantum theory.

Another option would be to argue that all of these wave functions can be treated on a par by the wave function realist. In recent paper, David Schroeren has argued that versions of WFR that employ different choices of representation differ only haecceitistically, and concludes

that the most attractive strategy for the wavefunction realist is to reconceive of their view as a *role-based* thesis that is not committed to a haecceitistic fact about which space is inhabited by the fundamental field, but rather as the thesis that *some* appropriately structured space plays this role. (Schroeren, 2002, 2)

9 Conclusion

Because standard quantum mechanics (both relativistic and non-relativistic) relies so heavily on background spacetime structure for its formulation, it

seems that a theory that would underwrite the conclusions about the structure of spacetime advocated by a Spicy version of WFR would have to depart fairly radically from standard quantum theory. That is, a Spicy version of WFR would also have to be Constructivist.

A project of this sort faces several challenges. One would be to motivate the departure from standard quantum theory. Despite some of the rhetoric employed by some proponents of WFR, the project cannot be one of exploring the metaphysical implications of standard quantum theory. Reasons for pursuing the project would have to include reasons for making the move from the Configuration Spaces on which standard quantum wave functions are defined to high-dimensional spaces with considerably less structure.

Carrying out this project would involve specifying what the structure of the fundamental space is meant to be. Does it come with a built-in measure of sizes of regions of the space? If so, what is the relation between this measure and, say, the usual Euclidean measures on configuration spaces, that are invoked in the definitions of standard quantum wave functions? Another question is what the wave functions invoked in WFR assign to points in the fundamental space. And the fundamental space will have to have whatever structure is needed to formulate the dynamical laws of evolution of the wave functions of WFR. What structure this is requires an answer to the question: what *are* the dynamical laws of the evolution of these wave functions?

Nothing in what has been said is meant to suggest that the difficulties involved in answering these questions are insurmountable. But they should be faced by proponents of WFR, and facing them requires acknowledging that we don't yet have even one version of wave function realism on the table. The project must be regarded as very much a work in progress, a research programme with the aim of constructing a theory that will conform to the metaphysical constraints imposed by the proponents of WFR.

10 Acknowledgments

This chapter has its origin in my comments, which mostly took the form of questions, at a symposium on Alyssa Ney's *The World in The Wave Function*, organized by Claudio Calosi, held online on June 15, 2021. I am grateful to the organizers of that symposium and to its participants, and in particular to Alyssa Ney, for her thoughtful responses to my comments. I am also grateful to discussants in an online presentation for the Lisbon Philosophy

of Physics Seminar (May 2023), and, of course at the *NOMOS* conference in Athens, June 13, 2023.

This chapter is the fruit of discussions, over many years, on the topic, with many people. I can't thank all by name, but I would like to single out David Wallace, Thomas De Saegher, and, above all, David Albert and Alyssa Ney, who have both been very patient with my continuing puzzlement about WFR.

11 Appendix: the plurality of wave-function representations

This appendix will be a bit more technical than the rest of the paper. It's here for those who want to understand the details of what I've said. It will be familiar to some readers, but, since much of this is glossed over in introductory quantum mechanics textbooks, it's worth rehearsing here.

11.1 Hilbert spaces from function spaces

Let $\langle \Gamma, \mathcal{G}, \mu \rangle$ be a measure space. That is, Γ is some non-empty set, \mathcal{G} is a σ -algebra of subsets of Γ , which are the ones we're going to assign measures to, and μ is a measure on \mathcal{G} , which means that it's a non-negative, countably additive set function.

Let $L^2(\Gamma, \mu)$ be the set of complex-valued functions on Γ that are square-integrable with respect to μ . This is obviously a vector space over the complex numbers. To construct a Hilbert space we first define an inner product

$$(f, g) = \int f^* g d\mu, \quad (8)$$

which defines a norm

$$\|f\| = \sqrt{(f, f)}. \quad (9)$$

Next, we define an equivalence relation on $L^2(\Gamma, \mu)$:

$$f \sim g \text{ iff } \|f - g\| = 0. \quad (10)$$

Equivalently: $f \sim g$ iff the set of all points on which the two functions differ has measure zero. We now define $\mathcal{L}^2(\Gamma, \mu)$ as the set whose elements are *equivalence classes*, under the relation \sim , of elements of $L^2(\Gamma, \mu)$.

Given the inner product (8) on $L^2(\Gamma, \mu)$, it's straightforward to define an inner product on $\mathcal{L}^2(\Gamma, \mu)$: given equivalence classes ψ, ϕ in $\mathcal{L}^2(\Gamma, \mu)$, pick some element f of ψ and some element g of ϕ , and apply (8). The space $\mathcal{L}^2(\Gamma, \mu)$, equipped with this inner product, is a Hilbert space.

Some might find it distasteful that we are dealing with equivalence classes of functions, rather than functions. If you don't like equivalence classes, then say we're still dealing with functions, but have agreed to read the equals sign in any equation that relates two elements of our function space, as the equivalence relation \sim . That works, too.

11.2 Unitarily equivalent representations

Suppose we have a quantum theory of some physical system. This involves a mapping from the dynamical variables of the system (e.g. position, momentum, spin, ...) to some set of operators \mathcal{A} , which are required to satisfy the canonical commutation relations. A representation of the quantum theory on a Hilbert space \mathcal{H} involves a mapping π from operators in \mathcal{A} to operators on \mathcal{H} that respects the algebraic structure of \mathcal{A} . Given such a representation, any vector ψ in \mathcal{H} defines a state of \mathcal{A} as follows: the expectation value of the result of an experiment yielding a value for the physical observable corresponding to the operator A is given by,¹⁵

$$\langle A \rangle_\psi = (\psi, \pi(A)\psi) / (\psi, \psi). \quad (11)$$

Given a representation π_1 of \mathcal{A} on a Hilbert space \mathcal{H}_1 , and a unitary mapping $U : \mathcal{H}_1 \rightarrow \mathcal{H}_2$ from \mathcal{H}_1 to another Hilbert space \mathcal{H}_2 , we can define a new representation π_2 , which maps operators in \mathcal{A} to operators on \mathcal{H}_2 , via,

$$\pi_2(A) = U\pi_1(A)U^{-1}. \quad (12)$$

It's easy to see that the state represented by the vector $U\psi$ in \mathcal{H}_2 is the same as the state represented by the vector ψ in \mathcal{H}_1 .

When two Hilbert-space representations of a quantum theory are related by some unitary U as in (12), they are said to be *unitarily equivalent*.

Now suppose that we are dealing with a system of N spinless particles, with Configuration Space $(\mathbb{E}^3)^N$. Pick some coordinatization of Configuration Space, assigning points in $(\mathbb{R}^3)^N$ to points in $(\mathbb{E}^3)^N$. Suppose we have a Configuration Space representation of our theory in terms of functions on $(\mathbb{R}^3)^N$, using some measure μ to define our inner product. A natural choice of measure is

$$d\mu = dx_1 dx_2 \dots dx_{3N}, \quad (13)$$

but nothing rides on this particular choice. Our Hilbert Space is thus $\mathcal{H}_1 = \mathcal{L}^2((\mathbb{E}^3)^N, \mu)$. In this representation, a dynamical variable that is a function $F(x_1, x_2, \dots, x_{3N})$ of the coordinates is represented by an operator \hat{F} on

¹⁵Note that any two vectors that are simply multiples of each other define the same state. We can, therefore, choose the norm of the vector we're using arbitrarily, and, for convenience, it's often handy to work with vectors of unit norm. But this is a matter of convenience only. (This footnote is for those who have read, in the metaphysics literature, that it's a requirement of quantum mechanics that wave functions be normalized, that is, that the integral of the square of the wave function over all the space on which it's defined be equal to one. This is simply not true.)

$\mathcal{L}^2((\mathbb{E}^3)^N, \mu)$ whose operation on a function ψ is given by,

$$(\hat{F}\psi)(x_1, x_2, \dots, x_{3N}) = F(x_1, x_2, \dots, x_{3N})\psi(x_1, x_2, \dots, x_{3N}). \quad (14)$$

We assume also that we have chosen appropriate operators $(\hat{p}_1, \hat{p}_2, \dots, \hat{p}_{3N})$ on $\mathcal{L}^2((\mathbb{E}^3)^N, \mu)$ to represent the momenta conjugate to the coordinates.

Now choose any function g on $(\mathbb{E}^3)^N$ that is measurable with respect to μ , and is nonzero everywhere. Define a new measure μ_g by,

$$d\mu_g = |g|^2 d\mu. \quad (15)$$

With this measure in hand, we have another Hilbert space \mathcal{H}_2 consisting of equivalence classes of functions that are square-integrable with respect to μ_g .

We can define a unitary mapping U_g from \mathcal{H}_1 to \mathcal{H}_2 via,

$$U_g \psi = \psi/g. \quad (16)$$

(The reader should pause and check that this is, indeed, unitary.)

If an observable A is represented by an operator $\pi_1(A)$ in the \mathcal{H}_1 -representation, in the \mathcal{H}_2 representation we are constructing it is represented by the operator

$$\pi_2(A) = U_g \pi_1(A) U_g^{-1} \quad (17)$$

This mapping leaves operators representing observables that are functions of coordinates unchanged: in both representations, these are multiplication operators, involving the same functions of coordinates. The momenta transform as,

$$\hat{p}_i \rightarrow U_g \hat{p}_i U_g^{-1} = \hat{p}_i - \left(\frac{i\hbar}{g} \frac{\partial g}{\partial x_i} \right) \quad (18)$$

We thus have two unitarily equivalent Configuration Space representations of our quantum theory: π_1 on $\mathcal{L}^2((\mathbb{E}^3)^N, \mu)$, and π_2 on $\mathcal{L}^2((\mathbb{E}^3)^N, \mu_g)$. A quantum state that is represented by a wave function ψ in the first representation is represented by a function $\phi = \psi/g$ in the second representation.

A special case is the familiar class of gauge transformations, which is obtained when g is a function that has unit magnitude everywhere, that is, when it is a function of the form,

$$g = e^{i\alpha}, \quad (19)$$

for some smooth real-value function α . This leaves the measure μ unchanged, and (18) becomes,

$$\hat{p}_i \rightarrow \hat{p}_i + \hbar \frac{\partial \alpha}{\partial x_i}. \quad (20)$$

So, how different can two Configuration Space wave functions that represent the same quantum state be? Take a quantum state that is represented by a function ψ in one representation. Pick any measurable function ϕ that you want, provided only that ϕ is nonzero wherever ψ is. Define g , for points where ψ (and hence ϕ) is nonzero, by

$$g = \psi/\phi. \quad (21)$$

For points (if there are any) at which ψ is zero, fill out the definition of g any way you like, in such a way that g is nonzero everywhere. Construct, as above, a wave-function representation on $\mathcal{L}^2((\mathbb{E}^3)^N, \mu_g)$. In this representation, unitarily equivalent to the first, the function ϕ represents the quantum state that is represented by ψ in our original representation. The new function ϕ can be any measurable function you like, as long as it's nonzero everywhere ψ is. There really is no such thing as *the* wave function that represents a quantum state!

As mentioned in the main text, the transformation U_g leaves the *ratios* of wave functions corresponding to two distinct states unchanged. That is, if we consider two wave functions ψ_1 , ψ_2 , and let $\phi_i = U_g \psi_i$ be the wave functions in the new representation,

$$\frac{\phi_1(\mathbf{x})}{\phi_2(\mathbf{x})} = \frac{\psi_1(\mathbf{x})}{\psi_2(\mathbf{x})}. \quad (22)$$

On the principle that what is physically significant in a mathematical representation of a physical state of affairs is what doesn't depend on our arbitrary choices, it is not the absolute value of a Configuration Space wave function that is physically significant, but the *ratios* of two Configuration Space wave functions representing distinct quantum states.

References

- Albert, D. Z. (1996). Elementary quantum metaphysics. In J. T. Cushing, A. Fine, and S. Goldstein (Eds.), *Bohmian Mechanics and Quantum Mechanics: An Appraisal*, 277–284. Kluwer.
- Albert, D. Z. (2013). Wave function realism. In Ney and Albert, eds. (2013), 52–57.
- Albert, D. Z. (2015). Quantum mechanics and everyday life. In *After Physics*, 124–143. Harvard University Press.
- Albert, D. Z. (2019). Preliminary considerations on the emergence of space and time. In A. Cordero, ed., *Philosophers Look at Quantum Mechanics*, 87–95. Springer.
- Bell, J. S. (1990). Indeterminism and nonlocality. Video of a lecture delivered at CERN, 22 January 1990. Available on YouTube at <https://www.youtube.com/watch?v=319HtG-VZCU>.
- Brown, H. R. and P. R. Holland (1999). The Galilean covariance of quantum mechanics in the case of external fields. *American Journal of Physics* 67, 204–214.
- Chen, E. K. (2017). Our fundamental physical space: An essay on the metaphysics of the wave function. *The Journal of Philosophy* 114, 333–365.
- Dasgupta, S., R. Dotan, and B. Weslake (eds.) (2021). *Current Controversies in Philosophy of Science*. Routledge.
- Goldstein, S., J. Taylor, R. Tumulka, and N. Zanghì (2005). Are all particles identical? *Journal of Physics A: Mathematical and General* 38(7), 1567–1576.
- Ismael, J. (2015). Quantum mechanics. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2015 ed.). Metaphysics Research Lab, Stanford University.
- Levy-Leblond, J.-M. (1963). Galilei group and nonrelativistic quantum mechanics. *Journal of Mathematical Physics* 4, 776–788.

- Lewis, P. J. (2004). Life in configuration space. *The British Journal for the Philosophy of Science* 55, 713–729.
- Loewer, B. (1996). Humean supervenience. *Philosophical Topics* 24, 101–127.
- Maudlin, T. (2013). The Nature of the Quantum State. In Ney and Albert, eds., (2013), 126–153.
- Monton, B. (2002). Wave function ontology. *Synthese* 130, 265–277.
- Myrvold, W. C. (2015). What is a wave function? *Synthese* 292, 3247–3274.
- Myrvold, W. C. (2019). How could relativity be anything other than physical? *Studies in History and Philosophy of Modern Physics* 67, 137–143.
- Myrvold, W. C. (2020a). On the status of quantum state realism. In S. French and J. Saatsi (Eds.), *Scientific Realism and the Quantum*, 219–251. Oxford: Oxford University Press. Available at <https://philsci-archive.pitt.edu/16655/>.
- Myrvold, W. C. (2020b). Subjectivists about quantum probabilities should be realists about quantum states. In M. Hemmo and O. Shenker (Eds.), *Quantum, Probability, Logic: The Work and Influence of Itamar Pitowsky*, 449–465. Springer Scientific. Available at <https://philsci-archive.pitt.edu/16656/>.
- Ney, A. (2021a). Separability, locality, and higher dimensions in quantum mechanics. In Dasgupta *et al.*, eds. (2021), 75–89.
- Ney, A. (2021b). *The World in the Wave Function: A Metaphysics for Quantum Physics*. Oxford University Press.
- Ney, A. and D. Z. Albert (Eds.) (2013). *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*. Oxford University Press.
- Reed, M. and B. Simon (1980). *Methods of Modern Mathematical Physics I. Functional Analysis* (Revised and Enlarged Edition). Academic Press.
- Schroeren, D. (2002). Wavefunction realism does not ‘privilege position’. *Synthese* 200, 27.

Wallace, D. (2021). Against wavefunction realism. In Dasgupta *et al.*, eds. (2021), 63–74.

Wittgenstein, L. (1953). *Philosophical Investigations*. Macmillan Publishing Co.