

1

Introduction

1.1 Scope of this book

Quantum mechanics is a difficult subject, and this book is intended to help the reader overcome the main difficulties in the way to understanding it. The first part of the book, Chs. 2–16, contains a systematic presentation of the basic principles of quantum theory, along with a number of examples which illustrate how these principles apply to particular quantum systems. The applications are, for the most part, limited to *toy models* whose simple structure allows one to see what is going on without using complicated mathematics or lengthy formulas. The principles themselves, however, are formulated in such a way that they can be applied to (almost) any nonrelativistic quantum system. In the second part of the book, Chs. 17–25, these principles are applied to quantum measurements and various quantum paradoxes, subjects which give rise to serious conceptual problems when they are not treated in a fully consistent manner.

The final chapters are of a somewhat different character. Chapter 26 on decoherence and the classical limit of quantum theory is a very sketchy introduction to these important topics along with some indication as to how the basic principles presented in the first part of the book can be used for understanding them. Chapter 27 on quantum theory and reality belongs to the interface between physics and philosophy and indicates why quantum theory is compatible with a real world whose existence is not dependent on what scientists think and believe, or the experiments they choose to carry out. The Bibliography contains references for those interested in further reading or in tracing the origin of some of the ideas presented in earlier chapters.

The remaining sections of this chapter provide a brief overview of the material in Chs. 2–25. While it may not be completely intelligible in advance of reading the actual material, the overview should nonetheless be of some assistance to readers who, like me, want to see something of the big picture before plunging into

the details. Section 1.2 concerns quantum systems at a single time, and Sec. 1.3 their time development. Sections 1.4 and 1.5 indicate what topics in mathematics are essential for understanding quantum theory, and where the relevant material is located in this book, in case the reader is not already familiar with it. Quantum reasoning as it is developed in the first sixteen chapters is surveyed in Sec. 1.6. Section 1.7 concerns quantum measurements, treated in Chs. 17 and 18. Finally, Sec. 1.8 indicates the motivation behind the chapters, 19–25, devoted to quantum paradoxes.

1.2 Quantum states and variables

Both classical and quantum mechanics describe how physical objects move as a function of time. However, they do this using rather different mathematical structures. In classical mechanics the *state* of a system at a given time is represented by a point in a *phase space*. For example, for a single particle moving in one dimension the phase space is the x, p plane consisting of pairs of numbers (x, p) representing the position and momentum. In quantum mechanics, on the other hand, the state of such a particle is given by a complex-valued *wave function* $\psi(x)$, and, as noted in Ch. 2, the collection of all possible wave functions is a complex linear vector space with an inner product, known as a *Hilbert space*.

The physical significance of wave functions is discussed in Ch. 2. Of particular importance is the fact that two wave functions $\phi(x)$ and $\psi(x)$ represent distinct physical states in a sense corresponding to distinct points in the classical phase space if and only if they are *orthogonal* in the sense that their inner product is zero. Otherwise $\phi(x)$ and $\psi(x)$ represent *incompatible* states of the quantum system (unless they are multiples of each other, in which case they represent the same state). Incompatible states cannot be compared with one another, and this relationship has no direct analog in classical physics. Understanding what incompatibility does and does not mean is essential if one is to have a clear grasp of the principles of quantum theory.

A quantum *property*, Ch. 4, is the analog of a collection of points in a classical phase space, and corresponds to a *subspace* of the quantum Hilbert space, or the *projector* onto this subspace. An example of a (classical or quantum) property is the statement that the energy E of a physical system lies within some specific range, $E_0 \leq E \leq E_1$. Classical properties can be subjected to various logical operations: negation, conjunction (AND), and disjunction (OR). The same is true of quantum properties as long as the projectors for the corresponding subspaces commute with each other. If they do not, the properties are incompatible in much the same way as nonorthogonal wave functions, a situation discussed in Sec. 4.6.

An orthonormal basis of a Hilbert space or, more generally, a decomposition of the identity as a sum of mutually commuting projectors constitutes a *sample space of mutually-exclusive possibilities*, one and only one of which can be a correct description of a quantum system at a given time. This is the quantum counterpart of a sample space in ordinary probability theory, as noted in Ch. 5, which discusses how probabilities can be assigned to quantum systems. An important difference between classical and quantum physics is that quantum sample spaces can be mutually incompatible, and probability distributions associated with incompatible spaces cannot be combined or compared in any meaningful way.

In classical mechanics a *physical variable*, such as energy or momentum, corresponds to a real-valued function defined on the phase space, whereas in quantum mechanics, as explained in Sec. 5.5, it is represented by a Hermitian operator. Such an operator can be thought of as a real-valued function defined on a particular sample space, or decomposition of the identity, but not on the entire Hilbert space. In particular, a quantum system can be said to have a value (or at least a precise value) of a physical variable represented by the operator F if and only if the quantum wave function is in an eigenstate of F , and in this case the eigenvalue is the value of the physical variable. Two physical variables whose operators do not commute correspond to incompatible sample spaces, and in general it is not possible to simultaneously assign values of both variables to a single quantum system.

1.3 Quantum dynamics

Both classical and quantum mechanics have *dynamical laws* which enable one to say something about the future (or past) state of a physical system if its state is known at a particular time. In classical mechanics the dynamical laws are *deterministic*: at any given time in the future there is a unique state which corresponds to a given initial state. As discussed in Ch. 7, the quantum analog of the deterministic dynamical law of classical mechanics is the (time-dependent) Schrödinger equation. Given some wave function ψ_0 at a time t_0 , integration of this equation leads to a unique wave function ψ_t at any other time t . At two times t and t' these uniquely defined wave functions are related by a unitary map or *time development operator* $T(t', t)$ on the Hilbert space. Consequently we say that integrating the Schrödinger equation leads to *unitary* time development.

However, quantum mechanics also allows for a *stochastic* or *probabilistic* time development, analogous to tossing a coin or rolling a die several times in a row. In order to describe this in a systematic way, one needs the concept of a *quantum history*, introduced in Ch. 8: a sequence of quantum *events* (wave functions or subspaces of the Hilbert space) at successive times. A collection of mutually

exclusive histories forms a sample space or *family* of histories, where each history is associated with a projector on a *history Hilbert space*.

The successive events of a history are, in general, not related to one another through the Schrödinger equation. However, the Schrödinger equation, or, equivalently, the time development operators $T(t', t)$, can be used to assign probabilities to the different histories belonging to a particular family. For histories involving only two times, an initial time and a single later time, probabilities can be assigned using the *Born rule*, as explained in Ch. 9. However, if three or more times are involved, the procedure is a bit more complicated, and probabilities can only be assigned in a consistent way when certain *consistency conditions* are satisfied, as explained in Ch. 10. When the consistency conditions hold, the corresponding sample space or event algebra is known as a *consistent family* of histories, or a *framework*. Checking consistency conditions is not a trivial task, but it is made easier by various rules and other considerations discussed in Ch. 11. Chapters 9, 10, 12, and 13 contain a number of simple examples which illustrate how the probability assignments in a consistent family lead to physically reasonable results when one pays attention to the requirement that stochastic time development must be described using a *single* consistent family or framework, and results from incompatible families, as defined in Sec. 10.4, are not combined.

1.4 Mathematics I. Linear algebra

Several branches of mathematics are important for quantum theory, but of these the most essential is *linear algebra*. It is the fundamental mathematical language of quantum mechanics in much the same way that calculus is the fundamental mathematical language of classical mechanics. One cannot even define essential quantum concepts without referring to the quantum Hilbert space, a complex linear vector space equipped with an inner product. Hence a good grasp of what quantum mechanics is all about, not to mention applying it to various physical problems, requires some familiarity with the properties of Hilbert spaces.

Unfortunately, the wave functions for even such a simple system as a quantum particle in one dimension form an *infinite-dimensional* Hilbert space, and the rules for dealing with such spaces with mathematical precision, found in books on functional analysis, are rather complicated and involve concepts, such as Lebesgue integrals, which fall outside the mathematical training of the majority of physicists. Fortunately, one does not have to learn functional analysis in order to understand the basic principles of quantum theory. The majority of the illustrations used in Chs. 2–16 are toy models with a finite-dimensional Hilbert space to which the usual rules of linear algebra apply without any qualification, and for these models there are no mathematical subtleties to add to the conceptual difficulties of

quantum theory. To be sure, mathematical simplicity is achieved at a certain cost, as toy models are even less “realistic” than the already artificial one-dimensional models one finds in textbooks. Nevertheless, they provide many useful insights into general quantum principles.

For the benefit of readers not already familiar with them, the concepts of linear algebra in finite-dimensional spaces which are most essential to quantum theory are summarized in Ch. 3, though some additional material is presented later: tensor products in Ch. 6 and unitary operators in Sec. 7.2. Dirac notation, in which elements of the Hilbert space are denoted by $|\psi\rangle$, and their duals by $\langle\psi|$, the inner product $\langle\phi|\psi\rangle$ is linear in the element on the right and antilinear in the one on the left, and matrix elements of an operator A take the form $\langle\phi|A|\psi\rangle$, is used throughout the book. Dirac notation is widely used and universally understood among quantum physicists, so any serious student of the subject will find learning it well-worthwhile. Anyone already familiar with linear algebra will have no trouble picking up the essentials of Dirac notation by glancing through Ch. 3.

It would be much too restrictive and also rather artificial to exclude from this book all references to quantum systems with an infinite-dimensional Hilbert space. As far as possible, quantum principles are stated in a form in which they apply to infinite- as well as to finite-dimensional spaces, or at least can be applied to the former given reasonable qualifications which mathematically sophisticated readers can fill in for themselves. Readers not in this category should simply follow the example of the majority of quantum physicists: go ahead and use the rules you learned for finite-dimensional spaces, and if you get into difficulty with an infinite-dimensional problem, go talk to an expert, or consult one of the books indicated in the bibliography (under the heading of Ch. 3).

1.5 Mathematics II. Calculus, probability theory

It is obvious that *calculus* plays an essential role in quantum mechanics; e.g., the inner product on a Hilbert space of wave functions is defined in terms of an integral, and the time-dependent Schrödinger equation is a partial differential equation. Indeed, the problem of constructing explicit solutions as a function of time to the Schrödinger equation is one of the things which makes quantum mechanics more difficult than classical mechanics. For example, describing the motion of a classical particle in one dimension in the absence of any forces is trivial, while the time development of a quantum wave packet is not at all simple.

Since this book focuses on conceptual rather than mathematical difficulties of quantum theory, considerable use is made of toy models with a simple discretized time dependence, as indicated in Sec. 7.4, and employed later in Chs. 9, 12, and 13. To obtain their unitary time development, one only needs to solve a simple

difference equation, and this can be done in closed form on the back of an envelope. Because there is no need for approximation methods or numerical solutions, these toy models can provide a lot of insight into the structure of quantum theory, and once one sees how to use them, they can be a valuable guide in discerning what are the really essential elements in the much more complicated mathematical structures needed in more realistic applications of quantum theory.

Probability theory plays an important role in discussions of the time development of quantum systems. However, the more sophisticated parts of this discipline, those that involve measure theory, are not essential for understanding basic quantum concepts, although they arise in various applications of quantum theory. In particular, when using toy models the simplest version of probability theory, based on a finite discrete sample space, is perfectly adequate. And once the basic strategy for using probabilities in quantum theory has been understood, there is no particular difficulty — or at least no greater difficulty than one encounters in classical physics — in extending it to probabilities of continuous variables, as in the case of $|\psi(x)|^2$ for a wave function $\psi(x)$.

In order to make this book self-contained, the main concepts of probability theory needed for quantum mechanics are summarized in Ch. 5, where it is shown how to apply them to a quantum system at a single time. Assigning probabilities to quantum histories is the subject of Chs. 9 and 10. It is important to note that the basic concepts of probability theory are the same in quantum mechanics as in other branches of physics; one does not need a new “quantum probability”. What distinguishes quantum from classical physics is the issue of choosing a suitable sample space with its associated event algebra. There are always many different ways of choosing a quantum sample space, and different sample spaces will often be incompatible, meaning that results cannot be combined or compared. However, in any single quantum sample space the ordinary rules for probabilistic reasoning are valid.

Probabilities in the quantum context are sometimes discussed in terms of a *density matrix*, a type of operator defined in Sec. 3.9. Although density matrices are not really essential for understanding the basic principles of quantum theory, they occur rather often in applications, and Ch. 15 discusses their physical significance and some of the ways in which they are used.

1.6 Quantum reasoning

The Hilbert space used in quantum mechanics is in certain respects quite different from a classical phase space, and this difference requires that one make some changes in classical habits of thought when reasoning about a quantum system. What is at stake becomes particularly clear when one considers the two-

dimensional Hilbert space of a spin-half particle, Sec. 4.6, for which it is easy to see that a straightforward use of ideas which work very well for a classical phase space will lead to contradictions. Thinking carefully about this example is worthwhile, for if one cannot understand the simplest of all quantum systems, one is not likely to make much progress with more complicated situations. One approach to the problem is to change the rules of ordinary (classical) logic, and this was the route taken by Birkhoff and von Neumann when they proposed a special quantum logic. However, their proposal has not been particularly fruitful for resolving the conceptual difficulties of quantum theory.

The alternative approach adopted in this book, starting in Sec. 4.6 and summarized in Ch. 16, leaves the ordinary rules of propositional logic unchanged, but imposes conditions on what constitutes a *meaningful* quantum description to which these rules can be applied. In particular, it is never meaningful to combine incompatible elements — be they wave functions, sample spaces, or consistent families — into a single description. This prohibition is embodied in the *single-framework* rule stated in Sec. 16.1, but already employed in various examples in earlier chapters.

Because so many mutually incompatible frameworks are available, the strategy used for describing the stochastic time development of a quantum system is quite different from that employed in classical mechanics. In the classical case, if one is given an initial state, it is only necessary to integrate the deterministic equations of motion in order to obtain a unique result at any later time. By contrast, an initial quantum state does not single out a particular framework, or sample space of stochastic histories, much less determine which history in the framework will actually occur. To understand how frameworks are chosen in the quantum case, and why, despite the multiplicity of possible frameworks, the theory still leads to consistent and coherent physical results, it is best to look at specific examples, of which a number will be found in Chs. 9, 10, 12, and 13.

Another aspect of incompatibility comes to light when one considers a tensor product of Hilbert spaces representing the subsystems of a composite system, or events at different times in the history of a single system. This is the notion of a *contextual* or *dependent* property or event. Chapter 14 is devoted to a systematic discussion of this topic, which also comes up in several of the quantum paradoxes considered in Chs. 20–25.

The basic principles of quantum reasoning are summarized in Ch. 16 and shown to be internally consistent. This chapter also contains a discussion of the intuitive significance of multiple incompatible frameworks, one of the most significant ways in which quantum theory differs from classical physics. If the principles stated in Ch. 16 seem rather abstract, readers should work through some of the examples found in earlier or later chapters or, better yet, work out some for themselves.

1.7 Quantum measurements

A quantum theory of measurements is a necessary part of any consistent way of understanding quantum theory for a fairly obvious reason. The phenomena which are specific to quantum theory, which lack any description in classical physics, have to do with the behavior of microscopic objects, the sorts of things which human beings cannot observe directly. Instead we must use carefully constructed instruments to amplify microscopic effects into macroscopic signals of the sort we can see with our eyes, or feed into our computers. Unless we understand how the apparatus works, we cannot interpret its macroscopic output in terms of the microscopic quantum phenomena we are interested in.

The situation is in some ways analogous to the problem faced by astronomers who depend upon powerful telescopes in order to study distant galaxies. If they did not understand how a telescope functions, cosmology would be reduced to pure speculation. There is, however, an important difference between the “telescope problem” of the astronomer and the “measurement problem” of the quantum physicist. No fundamental concepts from astronomy are needed in order to understand the operation of a telescope: the principles of optics are, fortunately, independent of the properties of the object which emits the light. But a piece of laboratory apparatus capable of amplifying quantum effects, such as a spark chamber, is itself composed of an enormous number of atoms, and nowadays we believe (and there is certainly no evidence to the contrary) that the behavior of aggregates of atoms as well as individual atoms is governed by quantum laws. Thus quantum measurements can, at least in principle, be analyzed using quantum theory. If for some reason such an analysis were impossible, it would indicate that quantum theory was wrong, or at least seriously defective.

Measurements as parts of gedanken experiments played a very important role in the early development of quantum theory. In particular, Bohr was able to meet many of Einstein’s objections to the new theory by pointing out that quantum principles had to be applied to the measuring apparatus itself, as well as to the particle or other microscopic system of interest. A little later the notion of measurement was incorporated as a fundamental principle in the standard interpretation of quantum mechanics, accepted by the majority of quantum physicists, where it served as a device for introducing stochastic time development into the theory. As von Neumann explained it, a system develops unitarily in time, in accordance with Schrödinger’s equation, until it interacts with some sort of measuring apparatus, at which point its wave function undergoes a “collapse” or “reduction” correlated with the outcome of the measurement.

However, employing measurements as a fundamental principle for interpreting quantum theory is not very satisfactory. Nowadays quantum mechanics is applied

to processes taking place at the centers of stars, to the decay of unstable particles in intergalactic space, and in many other situations which can scarcely be thought of as involving measurements. In addition, laboratory measurements are often of a sort in which the measured particle is either destroyed or else its properties are significantly altered by the measuring process, and the von Neumann scheme does not provide a satisfactory connection between the measurement outcome (e.g., a pointer position) and the corresponding property of the particle *before* the measurement took place. Numerous attempts have been made to construct a fully consistent measurement-based interpretation of quantum mechanics, thus far without success. Instead, this approach leads to a number of conceptual difficulties which constitute what specialists refer to as the “measurement problem.”

In this book all of the *fundamental* principles of quantum theory are developed, in Chs. 2–16, *without* making any reference to measurements, though measurements occur in some of the applications. Measurements are taken up in Chs. 17 and 18, and analyzed using the general principles of quantum mechanics introduced earlier. This includes such topics as how to describe a macroscopic measuring apparatus in quantum terms, the role of thermodynamic irreversibility in the measurement process, and what happens when two measurements are carried out in succession. The result is a consistent theory of quantum measurements based upon fundamental quantum principles, one which is able to reproduce all the results of the von Neumann approach and to go beyond it; e.g., by showing how the outcome of a measurement is correlated with some property of the measured system before the measurement took place.

Wave function collapse or reduction, discussed in Sec. 18.2, is not needed for a consistent quantum theory of measurement, as its role is taken over by a suitable use of conditional probabilities. To put the matter in a different way, wave function collapse is one method for computing conditional probabilities that can be obtained equally well using other methods. Various conceptual difficulties disappear when one realizes that collapse is something which takes place in the theoretical physicist’s notebook and not in the experimental physicist’s laboratory. In particular, there is no physical process taking place instantaneously over a long distance, in conflict with relativity theory.

1.8 Quantum paradoxes

A large number of quantum paradoxes have come to light since the modern form of quantum mechanics was first developed in the 1920s. A paradox is something which is contradictory, or contrary to common sense, but which seems to follow from accepted principles by ordinary logical rules. That is, it is something which ought to be true, but seemingly is not true. A scientific paradox may indicate that there is something wrong with the underlying scientific theory, which is quantum

mechanics in the case of interest to us. But a paradox can also be a prediction of the theory that, while rather surprising when one first hears it, is shown by further study or deeper analysis to reflect some genuine feature of the universe in which we live. For example, in relativity theory we learn that it is impossible for a signal to travel faster than the speed of light. This seems paradoxical in that one can imagine being on a rocket ship traveling at half the speed of light, and then shining a flashlight in the forwards direction. However, this (apparent) paradox can be satisfactorily explained by making consistent use of the principles of relativity theory, in particular those which govern transformations to moving coordinate systems.

A consistent understanding of quantum mechanics should make it possible to resolve quantum paradoxes by locating the points where they involve hidden assumptions or flawed reasoning, or by showing how the paradox embodies some genuine feature of the quantum world which is surprising from the perspective of classical physics. The formulation of quantum theory found in the first sixteen chapters of this book is employed in Chs. 20–25 to resolve a number of quantum paradoxes, including delayed choice, Kochen–Specker, EPR, and Hardy’s paradox, among others. (Schrödinger’s cat and the double-slit paradox, or at least their toy counterparts, are taken up earlier in the book, in Secs. 9.6 and 13.1, respectively, as part of the discussion of basic quantum principles.) Chapter 19 provides a brief introduction to these paradoxes along with two conceptual tools, quantum coins and quantum counterfactuals, which are needed for analyzing them.

In addition to demonstrating the overall consistency of quantum theory, there are at least three other reasons for devoting a substantial amount of space to these paradoxes. The first is that they provide useful and interesting examples of how to apply the basic principles of quantum mechanics. Second, various quantum paradoxes have been invoked in support of the claim that quantum theory is intrinsically *nonlocal* in the sense that there are mysterious influences which can, to take an example, instantly communicate the choice to carry out one measurement rather than another at point *A* to a distant point *B*, in a manner which contradicts the basic requirements of relativity theory. A careful analysis of these paradoxes shows, however, that the apparent contradictions arise from a failure to properly apply some principle of quantum reasoning in a purely local setting. Nonlocal influences are generated by logical mistakes, and when the latter are corrected, the ghosts of nonlocality vanish. Third, these paradoxes have sometimes been used to argue that the quantum world is not real, but is in some way created by human consciousness, or else that reality is a concept which only applies to the macroscopic domain immediately accessible to human experience. Resolving the paradoxes, in the sense of showing them to be in accord with consistent quantum principles, is thus a prelude to the discussion of quantum reality in Ch. 27.