Chapter 16

Quantum Reasoning

16.1 Some General Principles

There are some important differences between quantum and classical reasoning which reflect the different mathematical structure of the two theories. The most precise classical description of a mechanical system is provided by a point in the classical phase space, while the most precise quantum description is a ray or one-dimensional subspace of the Hilbert space. This in itself is not an important difference. What is more significant is the fact that two distinct points in a classical phase space represent mutually exclusive properties of the physical system: if one is a true description of the sytem, the other must be false. In quantum theory, on the other hand, properties are mutually exclusive in this sense only if the corresponding projectors are mutually orthogonal. Distinct rays in the Hilbert space need not be orthogonal to each other, and when they are not orthogonal, they do not correspond to mutually exclusive properties. As explained in Sec. 4.6, if the projectors corresponding to the two properties do not commute with one another, and are thus not orthogonal, the properties are (mutually) incompatible. The relationship of incompatibility means that the properties cannot be logically compared, a situation which does not arise in classical physics. The existence of this non-classical relationship of incompatibility is a direct consequence of assuming (following von Neumann) that the negation of a property corresponds to the orthogonal complement of the corresponding subspace of the Hilbert space; see the discussion in Sec. 4.6.

Quantum reasoning is (at least formally) identical to classical reasoning when using a single quantum *framework*, and for this reason it is important to be aware of the framework which is being used to construct a quantum description or carry out quantum reasoning. A framework is a Boolean algebra of commuting projectors based upon a suitable sample space, Sec. 5.2. The sample space is a collection of mutually orthogonal projectors which sum to the identity, and thus form a decomposition of the identity. A sample space of histories must also satisfy the consistency conditions discussed in Ch. 10.

In quantum theory there are always many possible frameworks which can be used to describe a given quantum system. While this situation can also arise in classical physics, as when one considers alternative coarse grainings of the phase space, it does not occur very often, and in any case classical frameworks are always mutually compatible, in the sense that they possess a common refinement. For reasons discussed in Sec. 16.4, compatible frameworks do not give rise to conceptual difficulties. By contrast, different quantum frameworks are generally incompatible, which means that the corresponding descriptions cannot be combined. As a consequence, when constructing a quantum description of a physical system it is necessary to restrict oneself to a single framework, or at least not mix results from incompatible frameworks. This *single framework rule* or *single family rule* has no counterpart in classical physics. Alternatively, one can say that in classical physics the single framework rule is always satisfied, for reasons indicated in Sec. 26.6, so one never needs to worry about it.

Quantum dynamics differs from classical Hamiltonian dynamics in that the latter is deterministic: given a point in phase space at some time, there is a unique trajectory in phase space representing the states of the system at earlier or later times. In the quantum case, the dynamics is stochastic: even given a precise state of the system at one time, various alternatives can occur at other times, and the theory only provides probabilities for these alternatives. (Only in the exceptional case of unitary histories, see Secs. 8.7 and 10.3, is there a unique (probability one) possibility at each time, and thus a deterministic dynamics.) Stochastic dynamics requires both the specification of an appropriate sample space or family of histories, as discussed in Ch. 8, and also a rule for assigning probabilities to histories. The latter, see Chs. 9 and 10, involves calculating weights for the histories using the unitary time development operators $T(t', t)$, equivalent to solving Schrödinger's equation, and then combining these with contingent data, typically an initial condition. Consequently, the reasoning process involved in applying the laws of quantum dynamics is somewhat different from that used for a deterministic classical system.

Probabilities can be consistently assigned to a family of histories of an isolated quantum system using the laws of quantum dynamics only if the family is represented by a Boolean algebra of projectors satisfying the *consistency conditions* discussed in Ch. 10. A family which satisfies these conditions is known as a *consistent family* or *framework*. Each framework has its own sample space, and the single framework rule says that the probabilities which apply to one framework cannot be used for a different framework, even for events or histories which are represented in both frameworks. It is, however, often possible to assign probabilities to elements of two or more distinct frameworks using the same initial data, as discussed below.

The laws of logic allow one to draw correct conclusions from some initial propositions, or "data", *assuming the latter are correct.* (Following the rules does not by itself always lead to the right answer; the principle of "garbage in, garbage out" was known to ancient logicians, though no doubt they worded it differently.) This is the sort of quantum reasoning with which we are concerned in this chapter. Given some facts or features of a quantum system, the "initial data", what else can we say about it? What conclusions can we draw by applying the principles of quantum theory? For example, an atom is in its ground state and a fast muon passes by 1 nm away: Will the atom be ionized? The "initial data" may simply be the initial state of the quantum system, but could also include information about what happens later, as in the specific example discussed in Sec. 16.2 below. Thus "initial" refers to what is given at the beginning of the logical argument, not necessarily some property of the quantum system which occurred before something else that one is interested in.

The first step in drawing conclusions from initial data consists in expressing the latter in proper quantum mechanical terms. In a typical situation the data are embedded in a sample space of mutually exclusive possibilities by assigning probabilities to the elements of this space. This includes the case in which the initial data identify a unique element of the sample space that is assigned a probability of 1, while all other elements have probability zero. If the initial data include information about the system at different times, the Hilbert space must, of course, be the Hilbert space of histories, and the sample space will consist of histories. See the example in Sec. 16.2 below. Initial data can also be expressed using a density matrix thought of as a pre-probability, see Sec. 15.6. Initial data which cannot be expressed in appropriate quantum terms cannot be used to initiate a quantum reasoning process, even if they make good classical sense.

Once the initial data have been embedded in a sample space, and probabilities have been assigned in accordance with quantum laws, *the reasoning process follows the usual rules of probability theory.* This means that, in general, the conclusion of the reasoning process will be a set of probabilities, rather than a definite result. However, if a consequence can be inferred with probability 1, we call it "true", while if some event or history has probability 0, it is "false", always assuming that the initial data are "true".

It is worth emphasizing once again that the peculiarities of quantum theory do not manifest themselves as long as one is using a *single* sample space and the corresponding event algebra. Instead, they come about because there are *many different* sample spaces in which one can embed the initial data. Hence the conclusions one can draw from those data depend upon which sample space is being used. This multiplicity of sample spaces poses some special problems for quantum reasoning, and these will be discussed in Secs. 16.3 and 16.4, after considering a specific example in the next section.

There are many other sorts of reasoning which go on when quantum theory is applied to a particular problem; e.g., the correct choice of boundary conditions for solving a differential equation, the appropriate approximation to be employed for calculating the time development, the use of symmetries, etc. These are not included in the present discussion because they are the same as in classical physics.

16.2 Example: Toy Beam Splitter

Consider the toy beam splitter with a detector in the c output channel shown in Fig. 12.2 on page 144 and discussed in Sec. 12.2. Suppose that the initial state at $t = 0$ is $|0a, 0\hat{c}\rangle$: the particle in the a entrance channel to the beam splitter, and the detector in its $0\hat{c}$ "ready" state. Also suppose that at $t = 3$ the detector is in its $1\hat{c}$ state indicating that the particle has been detected. These pieces of information about the the system at $t = 0$ and $t = 3$ constitute the initial data as that term was defined in Sec. 16.1. We shall also make use of a certain amount of "background" information: the structure of the toy model and its unitary time transformation, as found in Sec. 12.2.

In order to draw conclusions from the initial data, they must be embedded in an appropriate sample space. A useful approach is to begin with a relatively coarse sample space, and then refine it in different ways depending upon the sorts of questions one is interested in. One choice for the initial, coarse sample space is the set of histories

$$\begin{aligned} X^* &= [0a, 0\hat{c}] \odot I \odot I \odot [1\hat{c}], \\ X^\circ &= [0a, 0\hat{c}] \odot I \odot I \odot [0\hat{c}], \\ X^z &= \quad R \quad \odot I \odot I \odot \quad I \end{aligned} \tag{16.1}$$

for the times $t = 0, 1, 2, 3$, where $R = I - [0a, 0\hat{c}]$. Here the superscript $*$ stands for the triggered and $\circ$ for the ready state of the detector at $t = 3$. The sum of these projectors is the history identity $\breve{I}$, and it is easy to see that the consistency conditions are satisfied in view of the orthogonality of the initial and final states, Sec. 11.3. Since X^* is the only member of (16.1) consistent with the initial data, it is assigned probability 1, and the others are assigned probability 0.

Where was the particle at $t = 1$? The histories in (16.1) tell us nothing about any property of the system at $t = 1$, since the identity I is uninformative. Thus in order to answer this question we need to refine the sample space. This can be done by replacing X^* with the three history projectors

$$\begin{aligned} X^{*c} &= [0a, 0\hat{c}] \odot [1c] \odot I \odot [1\hat{c}], \\ X^{*d} &= [0a, 0\hat{c}] \odot [1d] \odot I \odot [1\hat{c}], \\ X^{*p} &= [0a, 0\hat{c}] \odot \;\; P \;\; \odot I \odot [1\hat{c}], \end{aligned} \tag{16.2}$$

whose sum is X^*, where

$$P = I - [1c] - [1d] \tag{16.3}$$

is the projector for the particle to be someplace other than sites $1c$ or $1d$. The weights of X^{*d} and X^{*p} are 0, given the dynamics as specified in Sec. 12.2. The history X° can be refined in a similar way, and the weights of $X^{\circ c}$ and $X^{\circ p}$ are 0. (We shall not bother to refine X^z, though this could also be done if one wanted to.) Consistency is easily checked.

When one refines a sample space, the probability associated with each of the elements of the original space is divided up among their replacements in proportion to their weights, as explained in Sec. 9.1. Consequently, in the refined sample space, X^{*c} has probability 1, and all the other histories have probability 0. Note that while X^{*d} and X^{*p} are consistent with the initial data, the fact that they have zero weight (are dynamically impossible) means that they have zero probability. From this we conclude that the initial data imply that the particle has the property $[1c]$, meaning that it is at the site $1c$, at $t = 1$. That is, $[1c]$ at $t = 1$ is true if one assumes the initial data are true.

Given the same initial data, one can ask a different question: At $t = 1$, was the particle in one or the other of the two states

$$|1\bar{a}\rangle = \big(|1c\rangle + |1d\rangle\big)/\sqrt{2}, \quad |1\bar{b}\rangle = \big(-|1c\rangle + |1d\rangle\big)/\sqrt{2} \tag{16.4}$$

resulting from the unitary evolution of $|0a\rangle$ and $|0b\rangle$ (see (12.2))? To answer this question, we use an alternative refinement of the sample space (16.1), in which X^* is replaced with the three histories

$$\begin{aligned} X^{*a} &= [0a, 0\hat{c}] \odot [1\bar{a}] \odot I \odot [1\hat{c}], \\ X^{*b} &= [0a, 0\hat{c}] \odot [1\bar{b}] \odot I \odot [1\hat{c}], \\ X^{*p} &= [0a, 0\hat{c}] \odot \;\; P \;\; \odot I \odot [1\hat{c}], \end{aligned} \tag{16.5}$$

with P again given by (16.3). (Note that $[1\bar{a}] + [1\bar{b}]$ is the same as $[1c] + [1d]$.) A similar refinement can be carried out for X°. Both X^{*b} and X^{*p} have zero weight, so the initial data imply that the history X^{*a} has probability 1. Consequently, we can conclude that the particle is in the superposition state $[1\bar{a}]$ with probability 1 at $t = 1$. That is, $[1\bar{a}]$ at $t = 1$ is true if one assumes the initial data are true.

However, the family which includes (16.5) is incompatible with the one which includes (16.2), as is obvious from the fact that $[1c]$ and $[1\bar{a}]$ do not commute with each other. Hence the probability 1 (true) conclusion obtained using one family cannot be combined with the probability 1 conclusion obtained using the other family. We cannot deduce from the initial data that at $t = 1$ the particle was in the state $[1c]$ and also in the state $[1\bar{a}]$, for this is quantum nonsense. Putting together results from two incompatible frameworks in this way violates the single-framework rule. So which is the *correct* family to use in order to work out the *real* state of the particle at $t = 1$: should one employ (16.2) or (16.5)? This is not a meaningful question in the context of quantum theory, for reasons which will be discussed in Sec. 16.4 below.

Now let us ask a third question based on the same initial data used previously. Where was the particle at $t = 2$: was it at $2c$ or at $2d$? The answer is obvious. All we need to do is to replace (16.2) with a different refinement

$$
\begin{aligned}
X^{*c'} &= [0a, 0\hat{c}] \odot I \odot [2c] \odot [1\hat{c}],\\
X^{*d'} &= [0a, 0\hat{c}] \odot I \odot [2d] \odot [1\hat{c}],\\
X^{*p'} &= [0a, 0\hat{c}] \odot I \odot \ P' \ \odot [1\hat{c}],
\end{aligned}
\tag{16.6}
$$

with $P' = I - [2c] - [2d]$. Since $X^{*c'}$ has probability 1, it is certain, given the initial data, that the particle was at $2c$ at $t = 2$.

The same answer can be obtained starting with the sample space which includes the histories in (16.2), and refining it to include the history

$$
X^{*cc} = [0a, 0\hat{c}] \odot [1c] \odot [2c] \odot [1\hat{c}], \tag{16.7}
$$

which has probability 1, along with additional histories with probability 0. In the same way, one could start with the sample space which includes the histories in (16.5), and refine it so that it contains

$$
X^{*ac} = [0a, 0\hat{c}] \odot [1\bar{a}] \odot [2c] \odot [1\hat{c}], \tag{16.8}
$$

whose probability (conditional upon the initial data) is 1, plus others whose probability is 0. It is obvious that the sample space containing (16.7) is incompatible with that containing (16.8), since these two history projectors do not commute with each other. Nonetheless, either family can be used to answer the question "Where is the particle at $t = 2$?", and both give precisely the same answer: the initial data imply that it is at $2c$, and not someplace else.

16.3 Internal Consistency of Quantum Reasoning

The example in Sec. 16.2 illustrates the principles of quantum reasoning introduced in Sec. 16.1. It also exhibits some important ways in which reasoning about quantum systems differs from what one is accustomed to in classical physics. In deterministic classical mechanics one is used to starting from some initial state and integrating the equations of motion to produce a trajectory in which at each time the system is described by a single point in its phase space. Given this trajectory one can answer any question of physical interest such as, for example, the time dependence of the kinetic energy.

In quantum theory one typically (unitary histories are an exception) uses a rather different strategy. Instead of starting with a single well-defined temporal development which can answer all questions, one has to start with the physical questions themselves and use these questions to generate an appropriate framework in which they make sense. Once this framework is specified, the principles of stochastic quantum dynamics can be brought to bear in order to supply answers, usually in the form of probabilities, to the questions one is interested in.

One cannot use a single framework to answer all possible questions about a quantum system, because answering one question will require the use of a framework that is incompatible with another framework needed to address some other question. But even a particular question can often be answered using more than one framework, as illustrated by the third (last) question in Sec. 16.2. This multiplicity of frameworks, along with the rule which requires that a quantum description, or the reasoning from initial data to a conclusion, use only a *single framework*, raises two somewhat different issues. The first issue is that of internal consistency: if many frameworks are available, will one get the same answer to the same question if one works it out in different frameworks? We shall show that this is, indeed, the case. The second issue, discussed in the next section, is the intuitive significance of the fact that alternative incompatible frameworks can be employed for one and the same quantum system.

The internal consistency of quantum reasoning can be shown in the following way. Assume that $\mathcal{F}_1, \mathcal{F}_2, \dots \mathcal{F}_n$ are different consistent families of histories, which may be incompatible with one another, each of which contains the initial data and the other events, or histories, that are needed to answer a particular physical question. Each framework is a set of projectors which forms a Boolean algebra, and one can define $\mathcal{F}$ to be their set-theoretic intersection:

$$\mathcal{F} = \mathcal{F}_1 \cap \mathcal{F}_2 \cap \cdots \mathcal{F}_n. \tag{16.9}$$

That is, a projector Y is in $\mathcal{F}$ if and only if it is also in each $\mathcal{F}_j$, for $1 \leq j \leq n$. It is straightforward to show that $\mathcal{F}$ is a Boolean algebra of commuting history projectors: It contains the history identity $\breve{I}$; if it contains a projector Y, then it also contains its negation $I - Y$; and if it contains Y and Y', then it also contains $YY' = Y'Y$. These assertions follow at once from the fact that they are true of each of the $\mathcal{F}_j$. Furthermore, the fact that each $\mathcal{F}_j$ is a consistent family means that $\mathcal{F}$ is consistent; one can use the criterion in (10.21).

Since each $\mathcal{F}_j$ contains the projectors needed to represent the initial data, along with those needed to express the conclusions one is interested in, the same is true of $\mathcal{F}$. Consequently, the task of assigning probabilities using the initial data together with the dynamical weights of the histories, and then using probabilistic arguments to reach certain conclusions, can be carried out in $\mathcal{F}$. But since it can be done in $\mathcal{F}$, it can also be done in an identical fashion in any of the $\mathcal{F}_j$, as the latter contains all the projectors of $\mathcal{F}$. Furthermore, any history in $\mathcal{F}$ will be assigned the same weight in $\mathcal{F}$ and in any $\mathcal{F}_j$, since the weight $W(Y)$ is defined directly in terms of the history projector Y using a formula, (10.11), that makes no reference to the family which contains the projector. Consequently, the conclusions one draws from initial data about physical properties or histories will be identical in all frameworks which contain the appropriate projectors.

This internal consistency is illustrated by the discussion of the third (last) question in Sec. 16.2: $\mathcal{F}$ is the family based on the sample space containing (16.6), and $\mathcal{F}_1$ and $\mathcal{F}_2$ are two mutually incompatible refinements containing the histories in (16.7) and (16.8), respectively. One can use either $\mathcal{F}_1$ or $\mathcal{F}_2$ to answer the question "Where is the particle at $t = 2$?", and the answer is the same.

As well as providing a proof of consistency, the preceding remarks suggest a certain strategy for carrying out quantum reasoning of the type we are concerned with: Use the smallest, or coarsest framework which contains both the initial data and the additional properties of interest in order to analyze the problem. Any other framework which can be used for the same purpose will be a refinement of the coarsest one, and will give the same answers, so there is no point in going to extra effort. If one has some specific initial data in mind, but wants to consider a variety of possible conclusions, some of which are incompatible with others, then start off with the coarsest framework $\mathcal{E}$ which contains all the initial data, and refine it in the different ways needed to draw different conclusions.

This was the strategy employed in Sec. 16.2, except that the coarsest sample space that contains the initial data X^* consists of the two projectors X^* and $\breve{I} - X^*$, whereas we used a sample space (16.1) containing three histories rather than just two. One reason for using X° and X^z in this case is that each has a straightforward physical interpretation, unlike their sum $\breve{I} - X^*$. The argument for consistency given above shows that there is no harm in using a more refined sample space as a starting point for further refinements, as long as it allows one to answer the questions one is interested in, for in the end one will always get precisely the same answer to any particular question.

16.4 Interpretation of Multiple Frameworks

The example of Sec. 16.2 illustrates a situation which arises rather often in reasoning about quantum systems. The initial data $\mathcal{D}$ can be used in various different frameworks $\mathcal{F}_1, \mathcal{F}_2 \ldots$, to yield different conclusions $\mathcal{C}_1, \mathcal{C}_2 \ldots$. The question then arises as to the relationship among these different conclusions. In particular, can one say that they all apply simultaneously to the same physical system? Generally the conclusions are expressed in terms of probabilities that are greater than 0 and less than 1, and thus involve some uncertainty. But sometimes, and we deliberately focused on this situation in the example in Sec. 16.2, one concludes that an event (or history) has probability 1, in which case it is natural to interpret this as meaning that the event actually occurs, or is a "true" consequence of the initial data. Similarly, probability 0 can be interpreted to mean that the event does not occur, or is "false".

If two or more frameworks are *compatible*, there is nothing problematical in supposing that the corresponding conclusions apply simultaneously to the same physical system. The reason is that compatibility implies the existence of a common refinement, a framework $\mathcal{G}$ which contains the projectors necessary to describe the initial data and all of the conclusions. The consistency of quantum reasoning, Sec. 16.3, means that the conclusions $\mathcal{C}_j$ will be identical in $\mathcal{F}_j$ and in $\mathcal{G}$. Consequently one can think of $\mathcal{F}_1, \mathcal{F}_2 \ldots$ as representing alternative "views" or "perspectives" of the same physical system, much as one can view an object, such as a teacup, from various different angles. Certain details are visible from one perspective and others from a different perspective, but there is no problem in supposing that they all from part of a single correct description, or that they are all simultaneously true, for the object in question.

In the example considered in Sec. 16.2, $\mathcal{F}_1$ could be the framework based on (16.2), which allows one to describe the position of the particle at $t = 1$, but not for any other $t > 0$, and $\mathcal{F}_2$ the one based on (16.6), which provides a description of the position of the particle at $t = 2$, but not at

$t = 1$. Their common refinement provides a description of the position of the particle at $t = 1$ and $t = 2$, and $\mathcal{F}_1$ and $\mathcal{F}_2$ can be thought of as supplying complementary parts of this description.

Conceptual difficulties arise, however, when two or more frameworks are *incompatible*. Again with reference to the example in Sec. 16.2, let $\mathcal{F}_3$ be the framework based on (16.5). It is incompatible with $\mathcal{F}_1$, because X^{*c} in (16.2) and X^{*a} in (16.5) do not commute with each other other, since the projectors $[1c]$ and $[1\bar{a}]$ at $t = 1$ do not commute. From the initial data one can conclude using $\mathcal{F}_1$ that the particle possesses the property $[1c]$ at $t = 1$ with probability 1. Using $\mathcal{F}_2$ and the same initial data, one concludes that the particle has the property $[1\bar{a}]$ at $t = 1$, again with probability 1. But even though $[1c]$ and $[1\bar{a}]$ are both "true" (probability 1) consequences of the initial data, one cannot think of them as representing properties of the particle which are simultaneously true in the same sense one is accustomed to when thinking about classical systems, for there is no property corresponding to $[1c]$ AND $[1\bar{a}]$, just as there is no property corresponding to $S_z = +1/2$ AND $S_x = +1/2$ for a spin-half particle.

The conceptual difficulty goes away if one supposes that the two incompatible frameworks are being used to describe two distinct physical systems that are described by the same initial data, or the same system during two different runs of an experiment. In the case of two separate but identical systems, each with Hilbert space $\mathcal{H}$, the combination is described by a tensor product $\mathcal{H} \otimes \mathcal{H}$, and employing $\mathcal{F}_1$ for the first and $\mathcal{F}_3$ for the second is formally the same as a single consistent family for the combination. This is analogous to the fact that while $S_z = +1/2$ AND $S_x = +1/2$ for a spin-half particle is quantum nonsense, there is no problem with the statement that $S_z = +1/2$ for one particle and $S_x = +1/2$ for a different particle. In the same way, different experimental runs for a single system must occur during different intervals of time, and the tensor product $\breve{\mathcal{H}} \odot \breve{\mathcal{H}}$ of two history Hilbert spaces plays the same role as $\mathcal{H} \otimes \mathcal{H}$ for two distinct systems.

Incompatible frameworks do give rise to conceptual problems when one tries to apply them to the *same* system during the *same* time interval. To be sure, there is never any harm in constructing as many alternative descriptions of a quantum system as one wants to, and writing them down on the same sheet of paper. The difficulty comes about when one wants to think of the results obtained using incompatible frameworks as all referring simultaneously to the same physical system, or tries to combine the results of reasoning based upon incompatible frameworks. It is this which is forbidden by the single framework rule of quantum reasoning.

Note, by the way, that in view of the internal consistency of quantum reasoning discussed in Sec. 16.3, it is never possible, even using incompatible frameworks, to derive *contradictory* results starting from the same initial data. Thus for the example in Sec. 16.2, the fact that there is a framework in which one can conclude with certainty that the particle is at the site $1c$ at $t = 1$ means there cannot be another framework in which one can conclude that the particle is someplace else at $t = 1$, or that it can be at site $1c$ with some probability less than 1. Any framework which contains both the initial data and the possibility of discussing whether the particle is or is not at the site $1c$ at $t = 1$ will lead to precisely the same conclusion as $\mathcal{F}_1$. This does not contradict the fact that in $\mathcal{F}_3$ the particle is predicted to be in a state $[1\bar{a}]$ at $t = 1$: $\mathcal{F}_3$ does not contain $[1c]$, and thus in this framework one cannot address the question of whether the particle is at the site $1c$ at $t = 1$.

Even though the single framework rule tells us that the result $[1c]$ from framework $\mathcal{F}_1$ and the result $[1\bar{a}]$ from $\mathcal{F}_3$ cannot be combined or compared, this state of affairs is intuitively rather troubling, for the following reason. In classical physics whenever one can draw the conclusion through one line of reasoning that a system has a property P, and through a different line of reasoning that it has the property Q, then it is correct to conclude that the system possesses both properties simultaneously. Thus if P is true (assuming the truth of some initial data) and Q is also true (using the same data), then it is always the case that P AND Q is true. By contrast, in the case we have been discussing, $[1c]$ is true (a correct conclusion from the data) in $\mathcal{F}_1$, $[1\bar{a}]$ is true if we use $\mathcal{F}_3$, while the combination $[1c]$ AND $[1\bar{a}]$ is not even meaningful as a quantum property, much less true!

When viewed from the perspective of quantum theory, see Ch. 26, classical physics is an approximation to quantum theory in certain circumstances in which the corresponding quantum description requires only a single framework (or, which amounts to the same thing, a collection of *compatible* frameworks). Thus the problem of developing rules for correct reasoning when one is confronted with a multiplicity of *incompatible* frameworks never arises in classical physics, or in our everyday "macroscopic" experience which classical physics describes so well. But this is precisely why the rules of reasoning which are perfectly adequate and quite successful in classical physics cannot be depended upon to provide reliable conceptual tools for thinking about the quantum domain. However deep-seated may be our intuitions about the meaning of "true" and "false" in the classical realm, these cannot be uncritically extended into quantum theory.

As probabilities can only be defined once a sample space has been specified, probabilistic reasoning in quantum theory necessarily depends upon the sample space and its associated framework. As a consequence, if "true" is to be identified with "probability 1", then the notion of "truth" in quantum theory, in the sense of deriving true conclusions from initial data that are assumed to be true, must necessarily depend upon the framework which one employs. This feature of quantum reasoning is sometimes regarded as unacceptable because it is hard to reconcile with an intuition based upon classical physics and ordinary everyday experience. But classical physics cannot be the arbiter for the rules of quantum reasoning. Instead, these rules must conform to the mathematical structure upon which quantum theory is based, and as has been pointed out repeatedly in previous chapters, this structure is significantly different from that of a classical phase space. To acquire a good "quantum intuition", one needs to work through various quantum examples in which a system can be studied using different incompatible frameworks. Several examples have been considered in previous chapters, and there are some more in later chapters. I myself have found the example of a beam splitter insider a box, Fig. 18.3 on page 219, particularly helpful. For additional comments on multiple incompatible frameworks, see Secs. 18.4 and 27.3.