# Contents

## CHAPTER 1

# Eight Experiments

Physics has traditionally been characterized as the science of matter in motion. Rough as this characterization is, it illuminates the standing of physics with respect to all other empirical sciences. Whatever else the objects of the various empirical sciences are, they are all instances of matter in motion. Every biological system, every economic system, every psychological system, every astronomical system is also matter in motion and so falls under the purview of physics. But not every physical system is biological or economic or psychological or astronomical. This is not to argue that these other empirical sciences reduce to physics, or that the other sciences do not provide an understanding of systems that is distinct from a purely physical account of them. Still, physics aspires to a sort of universality that is unique among empirical sciences and holds, in that sense, a foundational position among them.

The phrase "matter in motion" presents two targets for further analysis: "matter" and "motion." Present physics elucidates the "motion" of an object as its trajectory through space-time. A precise understanding of just what this is requires a precise account of the structure of space-time. The physical account of space-time structure has changed through the ages, and at present the best theory is the General Theory of Relativity. The nature of space-time itself and the geometrical structure of space-time is the topic of the companion volume to this one: *Philosophy of Physics: Space and Time*. The present volume addresses the question: What is matter? The best theory of matter presently available is quantum theory. Our main task is to understand just what quantum theory claims about the nature of the material constituents of the world.

As straightforward as this sounds, we must first confront a great paradox about modern physics. The two pillars on which modern physics rests are the General Theory of Relativity and quantum theory, but the status of these two theoretical systems is completely different. General Relativity is, in its own terms, completely clear and precise. It presents a novel account of space-time structure that takes some application and effort to completely grasp, but what the theory says is unambiguous. The more one works with it, the clearer it becomes, and there are no great debates among General Relativists about how to understand it. (The only bit of unclarity occurs exactly where one has to represent the distribution of matter in the theory, using the stress-energy tensor. Einstein remarked that that part of his theory is "low grade wood," while the part describing the space-time structure itself is "fine marble."[1]) In contrast, no consensus at all exists among physicists about how to understand quantum theory. There just is no precise, exact physical theory called "quantum theory" to be presented in these pages. Instead, there is raging controversy.

How can that be? After all, dozens and dozens of textbooks of quantum theory have been published, and thousands of physics students learn quantum theory every year. Some predictions of quantum theory have been subjected to the most exacting and rigorous tests in human history and have passed them. The whole microelectronics industry depends on quantum-mechanical calculations. How can the manifest and overwhelming empirical success of quantum theory be reconciled with complete uncertainty about what the theory claims about the nature of matter?

What is presented in the average physics textbook, what students learn and researchers use, turns out not to be a precise physical theory at all. It is rather a very effective and accurate *recipe* for making certain sorts of predictions. What physics students learn is how to use the recipe. For all practical purposes, when designing microchips and predicting the outcomes of experiments, this ability suffices. But if a physics student happens to be unsatisfied with just learning these mathematical techniques

[1] Einstein (1950), p. 84.

for making predictions and asks instead what the theory claims about the physical world, she or he is likely to be met with a canonical response: Shut up and calculate!

What about the recipe? Is it, at least, perfectly precise? It is not. John Stewart Bell pressed just this complaint:

> A preliminary account of these notions was entitled 'Quantum field theory without observers, or observables, or measurements, or systems, or apparatus, or wavefunction collapse, or anything like that'. That could suggest to some that the issue in question is a philosophical one. But I insist that my concern is strictly professional. I think that conventional formulations of quantum theory, and of quantum field theory in particular, are unprofessionally vague and ambiguous. Professional theoretical physicists ought to be able to do better.[2]

Bell's complaint is that the predictive recipe found in textbooks uses such terms as "observer" and "measurement" and "apparatus" that are not completely precise and clear. This complaint about quantum theory does not originate with Bell: Einstein famously asked whether a mouse could bring about drastic changes in the universe just by looking at it.[3] Einstein's point was that some formulations of quantum theory seek to associate a particular sudden change in the physical state of the universe ("collapse of the wavefunction") with acts of observation. If this is to count as a precise physical theory, then one needs a precise physical characterization of an observation. As Bell put it: "Was the wavefunction of the world waiting to jump for thousands of millions of years until a single-celled living creature appeared? Or did it have to wait a little longer, for a better qualified system . . . with a Ph.D.?"[4]

These imprecisions in the formulation of the quantum recipe do not have noticeable practical effects when it comes to making predictions. Physicists know well enough when a certain

---

[2] Bell (2004), p. 173.

[3] Reported by Hugh Everett in Everett (2012), p. 157.

[4] Bell (2004), p. 216.

laboratory operation is to count as an observation, and what it is an observation of. Quantum theory predicts the outcomes of these experiments with stunning accuracy. But if one's main interest is in the nature of the physical world rather than the pragmatics of generating predictions, this ability is of no solace. For the recipe simply does not contain any univocal account of the world itself. To illustrate this, the standard recipe does use a mathematical operation that can be called "collapse of the wavefunction." But if one asks whether that mathematical operation corresponds to a real physical change in the world itself, the recipe does not say. And practicing physicists do not agree on the answer. Some will refuse to hazard an opinion about it.

Bell's complaint might seem incredible. If the problems with quantum theory are not "merely philosophical" but rather consist of the theory being unprofessionally vague and ambiguous as physics, why don't the physics textbooks mention this? Much of the problem has been papered over by a misleading choice of terminology. A standard retort one might hear is this: Quantum mechanics as a physical theory is perfectly precise (after all, it has been used to make tremendously precise predictions!), but the interpretation of the theory is disputable. And, one might also hear, interpretation is a philosophical problem rather than a physical one. Physicists can renounce the desire to have any interpretation at all and just work with the theory. An interpretation, whatever it is, must be just an inessential luxury, like the heated seats in a car: It makes you feel more comfortable but plays no practical role in getting you from here to there.

This way of talking is misleading, because it does not correspond to what should be meant by a physical theory, or at least a fundamental physical theory. A physical theory should contain a physical *ontology*: What the theory postulates to exist as physically real. And it should also contain *dynamics*: laws (either deterministic or probabilistic) describing how these physically real entities behave. In a precise physical theory, both the ontology and the dynamics are represented in sharp mathematical terms. But it is exactly in this sense that the quantum-mechanical prediction-making recipe is not a physical theory. It does not specify what physically exists and how it behaves, but rather gives a (slightly vague) procedure for making statistical predictions about the outcomes of experiments. And what are often called "alternative interpretations of quantum theory" are rather alternative precise physical theories with exactly defined physical ontologies and dynamics that (if true) would explain why the quantum recipe works as well as it does.

Not every physical theory makes any pretense to provide a precisely characterized fundamental ontology. A physical theory may be put forward with the explicit warning that it is merely an approximation, that what it presents without further analysis is, nonetheless, derivative, and emerges from some deeper theory that we do not yet have in hand. In such a case, there may be circumstances in which the lowest level ontology actually mentioned by the theory is not precisely characterized. In the rest of this book, I will treat the theories under discussion as presenting a fundamental ontology that is not taken to be further analyzable, unless I indicate otherwise.

A precisely defined physical theory, in this sense, would never use terms like "observation," "measurement," "system," or "apparatus" in its fundamental postulates. It would instead say precisely *what exists* and *how it behaves*. If this description is correct, then the theory would account for the outcomes of all experiments, since experiments contain existing things that behave somehow. Applying such a physical theory to a laboratory situation would never require one to divide the laboratory up into "system" and "apparatus" or to make a judgment about whether an interaction should count as a measurement. Rather, the theory would postulate a physical description of the laboratory and use the dynamics to predict what the apparatus will (or might) do. Those predictions can then be compared to the data reported.

So far, then, we have distinguished three things: a physical theory, a recipe for making predictions, and the sort of data or phenomena that might be reported by an experimentalist. What is usually called "quantum theory" is a recipe or prescription, using some somewhat vague terms, for making predictions about data. If we are interested in the nature of the physical world, what we want is instead a theory—a precise articulation of what there is and how the physical world behaves, not just in the laboratory but at all places and times. The theory should be able to explain the success of the recipe and thereby also explain the phenomena.

Our order of investigation will start with some phenomena or data. We will try to report these phenomena in a "theory neutral" way, although in the end this will not quite be possible. But, as Aristotle said, any proper scientific investigation should start with what is clearer and more familiar to us and ascend to what is clearer by nature (*Physics* 184a16). We start with what we can see and try to end with an exactly articulated theory of what it really is.

Our phenomena are encapsulated in eight experiments.

## Experiment 1: The Cathode Ray Tube

The two ends of an electrical battery are called "electrodes." The positive electrode is the *anode*, and the negative one is the *cathode*. Run wires from these electrodes to two conductive plates, put an open aperture in the anode, place a phosphor-coated screen beyond the anode, and enclose the whole apparatus in an evacuated tube. Finally, add a controllable heating element to the cathode. This apparatus, minus the heating element, was invented by Ferdinand Braun in 1897 and later came to be called a *cathode ray tube* (CRT). The heating element was added in the 1920s by John B. Johnson and Harry Weinhart.

Our first experiment consists of adjusting the heating element so the cathode warms up. When the cathode is quite hot, a bright spot, roughly the shape of the aperture in the anode, appears on the phosphorescent screen (Figure 1a, 1b). As we turn the heating element down, the spot gets dimmer and dimmer. Eventually, the spot no longer shines steadily, but instead individual flashes of light appear in the same area (Figure 1c). As the heat is further lowered, these individual flashes become less and less frequent, eventually only appearing one at a time, with significant gaps between them. But if we keep track of these individual flashes, over time they trace out the same region as the original steady spot.

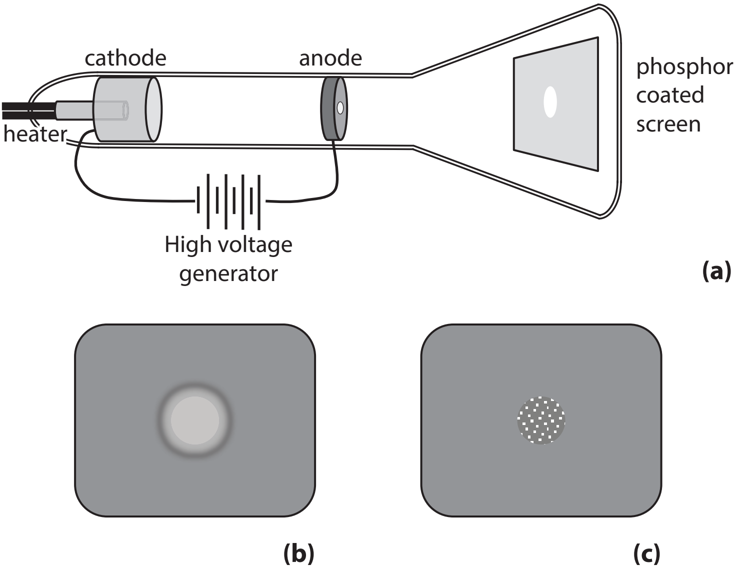

Figure 1

These are the phenomena or data. They immediately suggest certain hypotheses about what is going on inside the tube, but for the moment, we want to distinguish any such hypotheses from the data themselves. The phenomena suggest, for example, that something is going from the cathode (where the heating is applied), through the aperture in the anode, and to the phosphorescent screen. We can test this hypothesis by moving screen toward the anode while the spot is steady. The spot remains steady, and it narrows and brightens as it approaches the anode. Just in front of the anode, the spot is the same shape as the aperture. One can place a screen between the cathode and anode, where it will light with a larger, brighter, more diffuse glow. So there does seem to be something emitted from the cathode and going to the screen. Originally, this something was called *cathode rays*.

When we turn down the heat, these cathode rays exhibit a sort of discrete or grainy character, producing one flash at a time. We could not have predicted this behavior: The spot might have just dimmed uniformly without ever resolving into individual scintillations. These individual discrete events suggest a further hypothesis, namely, that the cathode rays are composed of a collection of individual particles. These hypothetical particles were eventually called *electrons*, and the whole cathode/anode apparatus is sometimes referred to as an *electron gun*.

The model suggested by the term "electron gun" is strengthened by the following fact. Increasing the voltage of the battery increases the "speed" of the cathode rays in the sense that if we measure how long it takes between connecting the battery and seeing the spot, it takes less time for higher voltages. This relation yields a narrative: Heating the cathode boils off electrons, which, being negatively charged, are repelled by the negatively charged cathode and attracted to the positively charged anode. The greater the voltage difference between the two, the faster the electrons will go, with some passing through the aperture in the anode and continuing on to the screen.

It is indeed difficult to resist this particle hypothesis, but for the moment, resist it we must. The postulation of individual particles that travel from the cathode to the screen is not itself part of the data, although it might be part of a theory meant to account for the data.

Skepticism about the physical existence of individual discrete particles in this experimental situation may seem excessively cautious, but our next two experiments point in another direction.

## Experiment 2: The Single Slit

If individual particles are flying from the cathode to the screen, then an object placed between the cathode and the screen might be able to affect these particles. As our first test of this hypothesis, we place a barrier with a single slit. The spot on the phosphorescent screen becomes long and thin, much as one might have anticipated (Figure 2a). Making the slit thinner in what we will call the $z$-direction initially makes the image thinner, again as one would expect. But beyond a certain point, a peculiar

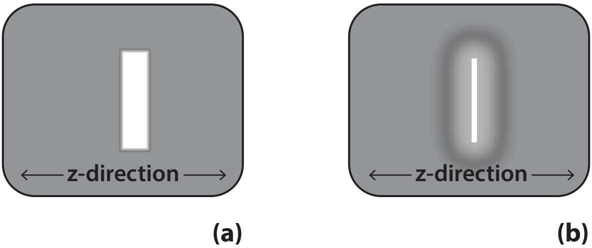

Figure 2

thing happens: making the slit even thinner results in the spot becoming wider and more spread out in the *z*-direction (Figure 2b). (In addition, the image starts to show some variation of brightness, with dark patches emerging. We leave that aside for now).

Our initial hypothesis of particles would not have hinted at this new development, but it is reminiscent of the familiar behavior of waves called *diffraction*. When a series of plane water waves hit a wide gap in a barrier (wide relative to the wavelength, i.e., distance between the crests), the wave train that gets through continues largely in the same direction, with just a little dissipation around the edges. But when it hits a very narrow gap, it creates a circular wave pattern on the other side that spreads farther upward and downward (Figures 3a and 3b). Crests are indicated by solid lines and troughs by dotted lines.

Since diffraction occurs when the size of the hole is small compared to the wavelength of the wave, the diffraction can be reduced by shortening the wavelength. And we find that the diffraction of our cathode rays is reduced as we increase the voltage between the cathode and the anode. So in this respect, our cathode rays behave somewhat like water waves, with the wavelength going down as the voltage goes up.

But it is still also the case that as we turn the heating element down, the glow goes from a steady state to a series of individual flashes, so in this sense, the phenomena suggest individual particles. The fact that the cathode rays (or electrons) produce phenomena associated with waves and also phenomena associated

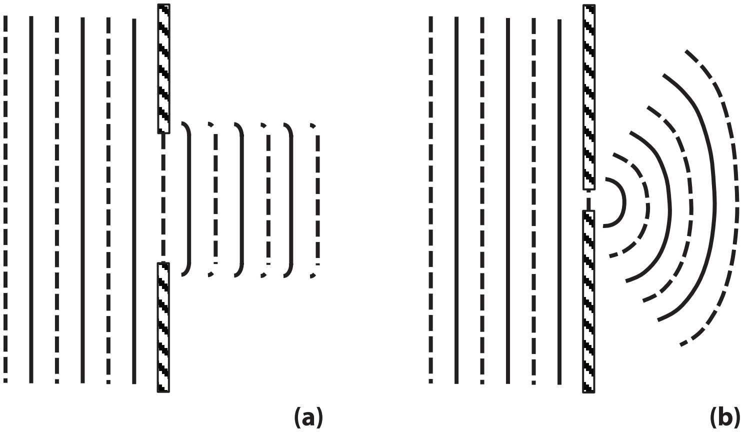

Figure 3

with particles is called *wave-particle duality*. But that is just a description of the phenomena, not an explanation of them.

## Experiment 3: The Double Slit

We are now in a position to describe the experiment most often associated with quantum theory: the double-slit experiment. In his classic *Lectures on Physics*, Richard Feynman is referring to the two-slit experiment when he says:

> We choose to examine a phenomenon which is impossible, *absolutely* impossible, to explain in any classical way, and which has in it the heart of quantum mechanics. In reality, it contains the *only* mystery. We cannot explain the mystery in the sense of "explaining" how it works. We will *tell* you how it works. In telling you how it works we will have told you about the basic peculiarities of all quantum mechanics.[5]

---

[5] Feynman, Leighton, and Sands (1975), Section 37-1.

Feynman is not correct when he says that there is no way to explain this phenomenon in a “classical” way, at least in one sense of “classical.” But this much is certainly true: the phenomenon is quite unexpected, and one does not really understand any physical theory that purports to account for the behavior of matter unless one understands how the theory accounts for this phenomenon. Belying Feynman’s pessimism, we will discuss several quite different exact physical theories, all of which can explain it.

Experiment 2 already demonstrates a behavior of cathode rays similar to that of water waves: diffraction. But an even more striking characteristic is associated with waves, namely, *interference*. Waves interfere because when they meet each other, they interact by *superposition*. For example, if the crest of one wave arrives at the same place as the equally deep trough of another, they cancel each other out; and if a crest meets a crest, they add to make a crest twice as tall. In Figures 3a,b, the solid lines represent the crests of the water waves, and the dotted lines represent the troughs. Now suppose instead of one hole or slit in the barrier we put two, and suppose that each slit is narrow enough to cause a lot of diffraction: in essence, each slit becomes the source of a set of circular wave patterns emanating from it (Figure 4). Where the crests of one wave meet the troughs of the other, they cancel out, and the water becomes still; where two crests meet or two troughs meet, the water is extremely agitated. Regions where crests coincide with crests and troughs with troughs are indicated by unbroken arrows, and regions where crests meet troughs by broken arrows. This superposition results in *interference bands* at the screen: regions of extremely high activity alternating with quiescent regions. Points on the screen where the difference in distances to the two slits is half a wavelength, or one and a half, or two and a half, and so forth (so the two arriving waves are out of phase) show little wave activity, and points where the difference is an integer number of wavelengths show lots of activity. Following the analogy with diffraction, then, we would expect alternating light and dark bands on the screen. This is indeed exactly what happens. Using a photographic emulsion and turning the heating element down yields a situation in which only individual dots

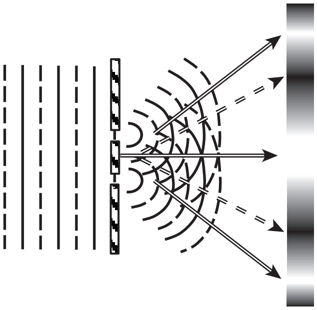

Figure 4

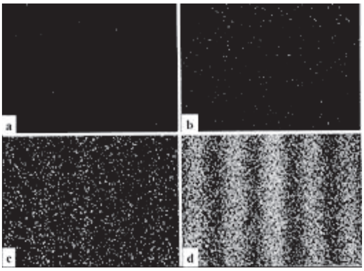

Figure 5. Credit: Reprinted courtesy of the Central Research Laboratory, Hitachi, Ltd., Japan.

appear on the screen, slowly accumulating to form the interference bands (Figure 5).

The two-slit interference experiment simultaneously displays properties we would naturally associate with particles (the individual discrete flashes or dots) and also properties we would naturally associate with waves (the interference bands). This is startling. But why would Feynman make the much stronger statement that the phenomenon cannot possibly be explained in any classical way?

Feynman's idea seems to be this: "to explain in a classical way" means to postulate the existence of individual particles that make their way, along one continuous path, from the cathode to the screen. Each such particle would therefore have to either pass through one slit or pass through the other (or loop around somehow to pass through both). Feynman calls the claim that each particle passes either through one slit or through the other "Proposition A." He then argues that Proposition A has some empirically testable consequences that turn out to be false, showing that we cannot accept it.

Suppose that each cathode ray that reaches the screen passes either just through the upper slit or just through the lower (leaving aside more rococo possibilities). Feynman reasons as follows. We can determine the final distribution of rays that pass through the upper slit by closing off the lower slit and seeing what happens. But we already know what happens: that is just Experiment 2. Similarly, we can close off the upper slit, in which case we get the same spread-out pattern just shifted over a bit. But the gaudily streaked interference pattern of the two-slit experiment is not the sum of these two experiments. Indeed, there are particular locations on the screen where spots will form if only the upper slit is open and spots will form if only the lower slit is open, *but no spots will form if both slits are open*.

Does it really follow from this observation, as Feynman suggests, that Proposition A cannot be true? In chapter 5, I will present a precise physical theory according to which each particle goes through exactly one slit and the interference bands only form when both slits are open. So there are ways to account for the data while validating Proposition A. What Feynman really seems to have in mind is not merely Proposition A, but also the additional proposition (call it Proposition B) that if an electron goes through one slit, then its later behavior will be the same regardless of whether the other slit is open. It is only with this second principle in place that one could infer that, given Proposition A, the distribution of flashes with both slits open would be the sum of the distributions with only one slit open. But Proposition B is not a proposition of classical physics, classical probability theory or classical logic. And the simple fact that there are locations on the screen where flashes occur if only one slit (whichever one) is open but never occur with both open already proves that *for each individual flash, the physical situation at the screen is sensitive to the condition of both slits*. This cannot be denied. What we want is a clear physical account of how it happens.

Denying Propositions A or B suggests that, in some sense, each electron or cathode ray interacts with both slits. And if this is true, then it is not surprising that the behavior at the screen can be sensitive to the fact that both slits are open. But for the electron or cathode ray to interact with both slits, it must somehow be spread out over a region large enough to encompass both slits, just as a water wave would have to be spread out that much. And in that case, the mystery is not so much how the behavior can depend on the state of both slits, but rather why the flashes on the screen occur at discrete, definite localizations. If a cathode ray really spreads out as a water wave does, how do the individual flashes manage to form?

As we will see, different precise theories embrace different horns of this dilemma. In some, the electron unproblematically goes through exactly one slit on its way to the screen, and the trick is to see how the other slit being open affects its later behavior. In others, the electron in some sense goes through both slits, and the trick is to account for the discrete flashes. But the situation is even more complicated: A slight modification of this experiment holds more surprises.

## Experiment 4: The Double Slit with Monitoring

Since our main puzzle concerns which slit, if either, the electron goes through on its trip from the anode to the screen, one might well ask: Why not just check? "Checking" means adding some new element to the experimental set-up designed to yield which way information about the electron, that is, information about which slit the particle went through. We will now explore a somewhat unrealistic and idealized modification of the experiment, but the effect of the modification on the phenomena is firmly based on quantum-mechanical principles.

With the thought that the electron is negatively charged, and that negatively charged particles attract positively charged ones, we might hit on the following scheme. Make a small, thin chamber in the screen between the two slits, and place a proton in a position exactly between the slits. Line the ends of the chamber with a substance that will emit a flash if a proton is absorbed (Figure 6).

If the electron goes through the upper slit, the proton will be attracted upward and the flash will occur at the top of the chamber, and if the electron goes through the lower slit, the flash will occur at the bottom. We can check the reliability of this monitoring mechanism by running it first with each slit closed. If it is 100% reliable, there will be a flash in the corresponding part of

Figure 6

the chamber when and only when there is a flash on the screen. That is, with the lower slit closed, there is a flash in the upper part of the chamber exactly when there is a flash on the screen, and similarly mutatis mutandis with the upper slit closed. We could also imagine less than perfect reliability: The corresponding flash might only happen 75% of the time, for example. We will consider this possibility presently.

Supposing we have achieved 100% reliability with only one slit open, what will happen when both slits are unblocked? Naively, we might expect to see the interference bands, as in Experiment 3, but now with additional information from the flash in the chamber about which slit the electron went through. Or, if the electron somehow goes through both slits and so would equally pull the proton up and down, maybe the proton will just remain symmetrically in the middle. Experiment 3 gives us no clue about the outcome.

As it turns out, this is what happens: Our completely reliable monitoring mechanism will continue to signal one slit or the other when there is a flash on the screen. And over time, about half of the electrons will be "seen" to have gone through the upper slit and about half through the lower. *But the interference bands will completely disappear*. The distribution of flashes on the screen will now be a simple sum of the distributions that occur when only one of the two slits is open. To put it somewhat poetically, when the path of the electron through the apparatus is observed, the behavior of the electrons changes from being wavelike (showing interference) to being particle-like (showing no interference). But notice that the "observer" in this poetic description is not even as sophisticated as a mouse. It is just a single proton whose own behavior has been coupled in the right way to that of the electron. There is something about that physical coupling that both destroys the interference and also seems to yield information about what the electron did.

What if we weaken the coupling between the electron and the proton? Suppose, for example, instead of reacting perfectly reliably when an electron goes by, the proton only moves from the central position 75% of the time (but always in the right direction, as checked with only one slit open)? What will we see then?

As the reliability of the monitor is reduced, the interference bands will slowly and continuously emerge. But as long as the behavior of the proton is correlated with the electron (with only one slit open) the interference bands will not be as strong as in Experiment 3. And the role played by the proton in destroying the interference is illustrated in a very striking way. If one divides the electron flashes on the screen into those that occur when the proton gives a result and those that occur when the proton stays in the central part of the chamber, the washed-out interference bands get split into two strikingly different sets. In the set where the proton indicates a slit, there is no interference at all, and in the set where it remains in the center, there is full interference. The total distribution is just the sum of these. As we progressively weaken the coupling with the proton, the interference bands progressively reemerge to full force.

One might well wonder how any clear and precise physical account of what is going on could yield this sort of behavior. What sort of pattern appears on the screen seems to depend on whether, in some sense, anyone or anything is "watching" the electron. But must the physical theory therefore define "watching" in order to be articulated? How can that be done? Has the observer somehow claimed a central place in physics? Many physicists over the years have drawn just this conclusion. Experiment 4 gives us some indication of the phenomena that led them to it. But the very simplicity of the "watcher" in this experiment is promising. There is little prospect of producing an exact physical characterization of something as large and complicated as a mouse. But a single proton, coupled by electrical attraction to an electron, is exactly the sort of thing we expect an exact physical theory to treat with complete precision. So Experiment 4 ought to give us some hope.

## EXPERIMENT 5: SPIN

Our previous experiments have illustrated some of the peculiarities of quantum theory. It is easy to see in these phenomena wave-particle duality, since individual flashes are suggestive of particles, and the collective interference patterns are suggestive of waves. Our last experiment illustrates how the physical role of observation might appear as a central theme. Even the simple single-slit diffraction experiment provides an instance of the famous Heisenberg uncertainty relations. Werner Heisenberg noticed that as our predictive abilities become better in some ways, they simultaneously become worse in others. The sorts of predictions that trade off in this way are called "complementary." One standard example of complementarity is position and momentum in a given direction. Narrowing the slit in Experiment 2 decreases uncertainty about where in the $z$-direction a particle that passes the slit will show up just beyond the slit, but it simultaneously increases uncertainty about its $z$-momentum (i.e., how fast and in what direction it is moving in the $z$-direction) at that point. This increased spread in possible $z$-momenta results in the widening of the image in the $z$-direction far from the slit. But so far we have not had much indication of the “quantum” in quantum theory. It is popularly thought that in quantum theory everything is quantized into discrete units. But in our examples so far, that is not so. Our cathode rays can appear as flashes at any location on the screen, for example.

The simplest physical property that exhibits quantization is called “spin,” and manifests itself as an intrinsic angular momentum of a particle. In classical physics, a spinning charged particle has a magnetic polarization. If an object has a north and south magnetic pole, then it will be deflected when travelling through an inhomogeneous magnetic field. Figure 7 shows a diagram of a Stern-Gerlach apparatus that produces this effect.

The apparatus is just a magnet, but because of the asymmetric geometry, the north pole creates a locally stronger magnetic field than does the south pole. A bar magnet in the field oriented with its north pole up and its south pole down would be pushed down, since the north will be repelled by the upper field more strongly than the south is repelled by the lower. Similarly, a bar magnet oriented the opposite way will move up, since the attraction of its south pole upward will overbalance the attraction of the north pole downward. A horizontally oriented bar magnet will not be pushed or pulled either way.

Our electrons are negatively charged, but it does not immediately follow that they have magnetic moments. Classically, a spinning electric charge does create a magnetic field. Hence an intrinsic magnetic moment of a particle is associated with its “spin,” irrespective of whether it originates in the actual spinning motion of anything. One way to check for such a magnetic moment is to pass a particle through a Stern-Gerlach apparatus to see whether it is deflected.

If one does this sort of experiment on our cathode rays, the outcome is somewhat unexpected.[6] Every electron is deflected either up or down, with none going straight through. Furthermore,

---

[6] The physics here is somewhat idealized, although again the basic principles are correct. In practice, this experiment was first done with silver atoms.

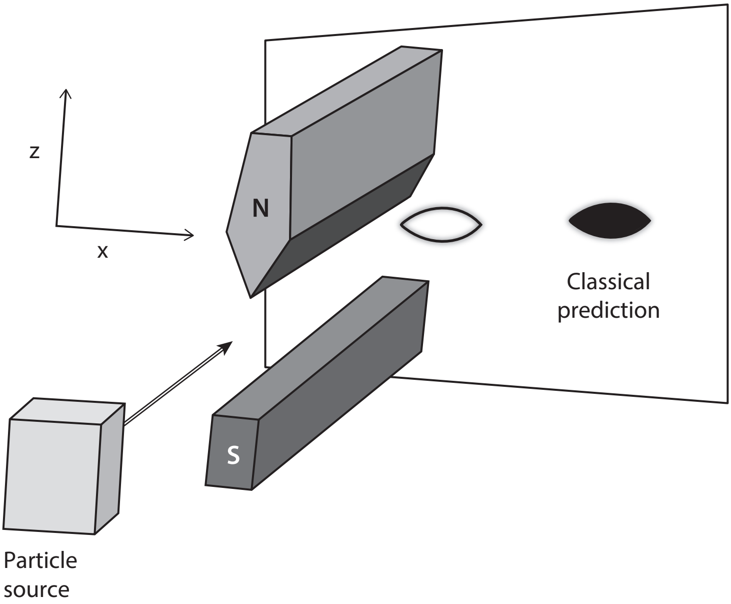

Figure 7

the amount of deflection is exactly the same in both cases. Our electron beam divides into two equally strong parts, one deflected toward the north pole of the apparatus (up-spin) and the other toward the south pole (down-spin). Or, more precisely, our steady lighted patch on the screen splits into two equally bright patches, one above and the other below the midline. And as we turn the beam intensity down, we again get individual flashes, about half in a small region in the upper part of the screen, half in an equally small region in the lower part. Particles in a beam that splits into exactly two parts are called “spin-1/2” particles. If it were to split into three parts, one going straight through, the particles would be spin-1 particles. A beam of spin-3/2 particles splits into four parts, and so on.

In Figure 7, the image on the screen looks like an eye, because the electron beam out at the edges does not go through an inhomogeneous field and so travels straight through. Stern and Gerlach’s

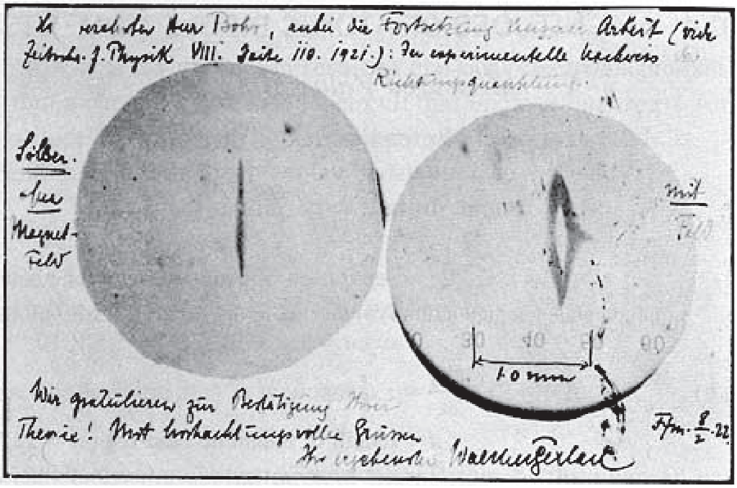

Hochverehrter Herr Bohr, anbei die Fortsetzung unserer Arbeit (vergl. Zeitschr. f. Physik VIII. Seite 110. 1921): Den experimentellen Nachweis der Richtungsquantelung.

Silber. Ohne Magnet-Feld

mit Feld

Wir gratulieren zur Bestätigung Ihrer Theorie! Mit hochachtungsvollen Grüssen Ihr ergebenster Walther Gerlach

Ffm. 8/2. 22

Figure 8. With permission of Niels Bohr Archive, Copenhagen.

actual data are shown in Figure 8 on a postcard that Gerlach sent to Niels Bohr. In our schematic diagrams, we will cut off the sides, so only the most separated parts of the beam are indicated.

The Stern-Gerlach apparatus is itself oriented in some spatial direction. Figure 7 designates the vertical direction as the *z*-direction and the horizontal one as the *x*-direction. If we twist the apparatus from the *z*-orientation to the *x*-orientation, the beam comes to split in the *x*-direction, as in the postcard. The apparatus can be set to have any spatial direction.

Since about half of the particles are deflected up and half deflected down, one naturally wonders whether some feature of each individual particle determines which way it goes. It not obvious how to resolve this question experimentally, but some additional experimental configurations are clearly relevant. Let a first Stern-Gerlach apparatus be oriented in the *z*-direction, splitting the beam into an upper and lower branch. Then place a second apparatus, also oriented in the *z*-direction, in each of these beams (Figure 9a). We might expect each beam to split again, but it does not: the whole upper beam is deflected up and the whole lower

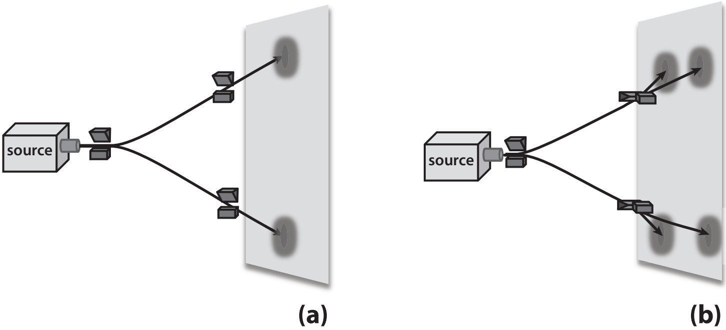

Figure 9

beam deflected down. So electron beams can be prepared so that each electron in them is disposed to be deflected in a particular way by a *z*-oriented apparatus. Those disposed to be deflected upward are called "*z*-spin up" electrons, and those disposed to be deflected down called "*z*-spin down."

We can also follow a *z*-oriented apparatus with *x*-oriented ones (Figure 9b). In this case, each beam splits 50-50. So preparing a beam so it can be predicted which way each electron will be deflected in the *z*-direction results in complete uncertainty about how it will be deflected in the *x*-direction. And testing the output of the *x*-apparatus with yet another *z*-apparatus reveals that the original preparation has been lost: the beam once again splits 50-50 up and down.

All electrons in a *z*-spin up beam get deflected up in the *z*-direction and only half do in the *x*-direction. What if we slowly rotate the second apparatus from the *z*-orientation to the *x*-orientation? Unsurprisingly, the proportion deflected in the up direction (with respect to the apparatus) varies smoothly from 1 to 0.5. More quantitatively, the proportion deflected up at the second apparatus is $\cos^2(\theta/2)$, where $\theta$ is the angle between the orientations of the two apparatuses.

Our spin experiments illustrate the quantization of spin, since each electron responds to the experimental condition in one of two possible ways. They also illustrate the Heisenberg uncertainty relations: The more certain it is how an electron will react to a *z*-oriented apparatus, the less certain it will be how the electron reacts to one that is *x*-oriented and vice versa. This is analogous to the situation with *z*-position and *z*-momentum in the single-slit experiment.

The quantization of spin offers particularly sharp and clear experimental possibilities that are the subject of our next laboratory configuration.

### EXPERIMENT 6: THE INTERFEROMETER

Our next experiment refines some of the phenomena we have already discussed. We have seen how a Stern-Gerlach apparatus can split an incoming beam of spin-1/2 particles into two beams. Those beams can be further manipulated and recombined in an experimental configuration that was originally developed for light by Ludwig Mach and Ludwig Zehnder, and hence is known as the Mach-Zehnder interferometer.

The first experiment is a slight variation on a spin experiment we have already discussed. Prepare an *x*-spin up beam of electrons and pass it through a *z*-oriented Stern-Gerlach device. We have already remarked that if we pass either of the output beams through an *x*-oriented apparatus, the beam will again split: apparently the *z*-oriented magnet "scrambles" the information about the prepared *x*-spin. In itself, this is not terribly surprising. The interaction of the beam with the new magnetic field could have all sorts of disruptive effects. But the Mach-Zehnder configuration allows us to steer the output beams of the *z*-oriented device back together, having been widely separated from each other for some time (Figure 10). A natural train of thought runs as follows: The *x*-spin of each separate output beam of the *z*-oriented magnet has been scrambled, with each particle equally likely to be deflected up or down. When two such scrambled beams are combined, the result should be just as scrambled. So the recombined beam should also be equally split if passed through an *x*-oriented magnet at point A in the figure.

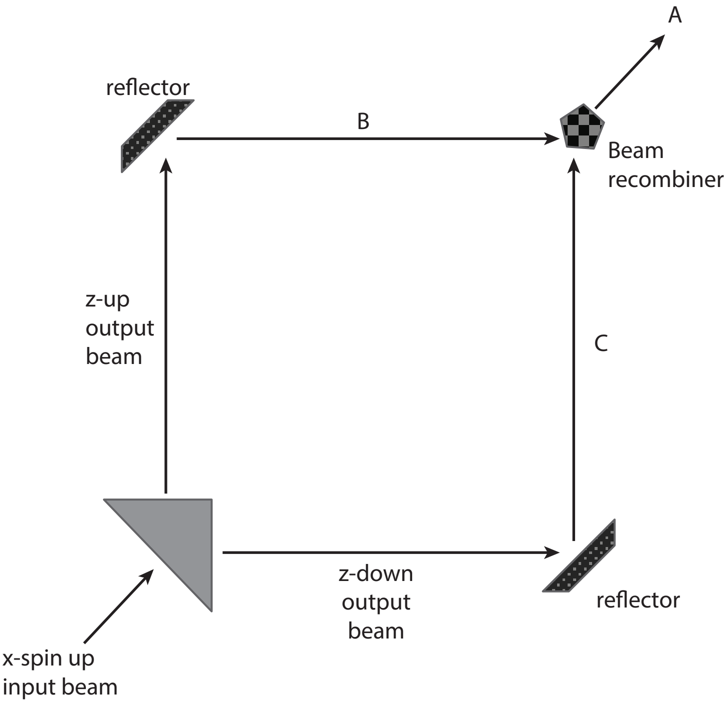

Figure 10

This, however, is not what happens. Every single electron is deflected up at point A even though half would have been deflected down if the beam had instead been checked by a pair of devices located at points B and C. And if *x*-spin down electrons are fed into the interferometer, we get a similar result: Half will be deflected down if the *x*-spin is checked at points B and C, but all are deflected down if the *x*-spin of the recombined beam is checked at A. Information about how the original beam was prepared is somehow transmitted through the splitting and recombination, even though that very information appears to have been lost half way through! And once again, these statistics hold even if we turn down the intensity of the incoming beam so only one electron goes through the interferometer at a time.

The interferometer set-up allows us a new opportunity: to intervene on one branch of the split beam before it is recombined. One intervention is particularly striking. We have already seen that the spin of an electron is like an intrinsic magnetic moment, akin to that of a classical spinning charge. In a classical theory, applying a magnetic field to such an object will cause it to precess (i.e., to slowly rotate in space). One sort of magnetic field, applied for the right amount of time, would cause the object to rotate though a full 360° and hence (apparently) return to the same state it started in. We can apply such a magnetic field to our electrons,[7] and check that the precession does return the beam to its initial state: an *x*-spin up beam remains *x*-spin up when the magnetic field is applied, an *x*-spin down beam remains *x*-spin down, *z*-spin up remains *z*-spin up, etc. Further, an *x*-spin up beam gets converted to *x*-down if the magnetic field is applied for half the time, just as one would expect if it were rotated through 180°. A device that applies the magnetic field for the full time is an example of what David Albert calls a "total of nothing box" because the observable statistics of any beam are unchanged by the application of the magnetic field.[8] As far as predictions are concerned, a beam that has had the magnetic field applied behaves just like one that has not. The foregoing remarks hold so long as the whole beam is subjected to the magnetic field.

But suppose we split the beam in the interferometer and apply the magnetic field to only one part of the split beam (at point B, for example) and then recombine the beams. This intervention has a dramatic effect on the outcome. Without the magnetic field, as we have seen, if we feed an *x*-spin up beam in, we get an *x*-spin up beam out after the recombination. But with the magnetic field in place, when we feed an *x*-spin up beam in, we get an *x*-spin *down* beam out. Every single electron is deflected down by an *x*-oriented magnet at the end, while without the magnetic field, every single electron is deflected up. In other words, every

---

[7] I am fudging the actual physics a bit: The experiment described here was carried out on neutrons rather than electrons. Neutrons also are spin-1/2 particles.

[8] Albert (1992), p. 11.

electron fed through our device is demonstrably sensitive to the physical conditions along both paths in the interferometer: A certain magnetic field applied either at point B or at point C (but not both) will alter the behavior of every electron that passes through.

This is not, strictly speaking, a new sort of observation. We have already seen in the two-slit interference that every electron is sensitive to the state of both slits. The Mach-Zehnder configuration brings out this fact in a particularly striking way, since the two paths through the interferometer can be made to diverge from each other by arbitrary distances. Nonetheless, an intervention on either branch can have an effect on every single electron.

## Experiment 7: The EPR Experiment

Unlike the interferometer, our final two experiments bring in fundamentally new features of quantum theory. Indeed, we are starting on the path to the most puzzling and astonishing physical phenomena predicted by the quantum formalism and verified in the laboratory. These phenomena essentially involve collections of particles rather than single particles or beams of single particles. So far, only Experiment 4, the Double Slit with Monitoring, has required more than one particle at a time. Experiment 4 demands this because the monitoring proton and passing electron must interact for the monitoring to occur. We now embark on a deeper investigation into such interactions, and into the information that the behavior of one particle can yield about another.

The first experiment is a modification, proposed by David Bohm, of an experimental situation described by Albert Einstein, Boris Podolsky, and Nathan Rosen in "Can Quantum-Mechanical Description of Reality Be Considered Complete?" (1935), now known as the EPR paper.[9] In that paper, the discussion concerned the positions and momenta of a pair of particles prepared in a special state. Bohm changed the example to use spin in

---

[9] Einstein, Podolsky, and Rosen (1935).

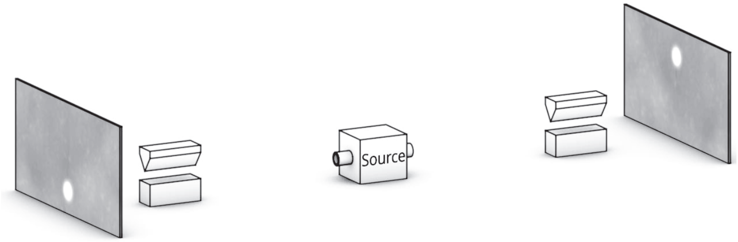

Figure 11

various directions rather than position and momentum, and we will follow his simpler example.

Unlike all the experiments described so far, the basic phenomena in the EPR experiment seem unremarkable. A pair of electrons is prepared in a particular quantum-mechanical state (called the “singlet” state) and is allowed to separate to an arbitrary distance from each other. Each electron is then passed through a Stern-Gerlach apparatus oriented in a specific spatial direction (Figure 11). For example, both of the devices might be oriented in the $z$-direction. In this case, the two electrons always behave in opposite ways: If one is deflected upward in the apparatus, the other is deflected downward. Therefore, by observing how one electron behaves, one can predict with perfect accuracy how the other will (or has).

So what is so remarkable about this? It is true that the behavior of one electron provides information about how the other will behave, but everyday instances of these sorts of correlations are commonplace. John Bell used the amusing example of the physicist Reinhold Bertlmann, who always wore socks of different colors.[10] Given this somewhat idiosyncratic choice of how to get dressed, the color of one sock (pink, say) provides information about the color of the other (not-pink).

---

[10] “Bertlmann’s Socks and the Nature of Reality,” reprinted as Chapter 16 in Bell (2004).

But notably, this prosaic account of the phenomenon essentially presupposes that the socks have their colors all along, from the time Bertlmann got dressed. If, somehow, neither sock had any definite color in the morning, if a sock only acquired a definite color some time later (when observed, for example), then the situation would be truly remarkable. It would be remarkable first because of the no-definite-color to definite-color transition. One would rightly wonder about the physics of that change. But even granting that, there is a residual surprise, for not only would the one sock have to somehow come to become actually pink at some point, but the other sock (which might be miles away) would also somehow have to become some color other than pink, so that the colors would always be different. This idea, that interacting with one sock can somehow not merely provide information about the other but actually affect the physical state of the other, is an example of the possibility of *quantum nonlocality*.

Einstein, Podolsky, and Rosen never took the possibility of such a nonlocal physical interaction between the socks (or the electrons) seriously. In fact, they thought the idea so absurd that they never imagined anyone would entertain it. What the EPR article pointed out was that to avoid such a strange "spooky action-at-a-distance" (in Einstein's famous phrase), one has to postulate that the two electrons described above have definite dispositions concerning how they would react to the magnets from the moment they are produced and separate from each other. One of the electrons has to be $z$-spin up and the other $z$-spin down from the outset. Otherwise, how could either be sensitive to the behavior of the other in the right way to preserve the perfect anticorrelation?

It is worthwhile to belabor this point a little. Imagine, as an analogy, that you and a friend are going to be subjected to the following ordeal. You are going to be taken into separate rooms and asked a yes-no question. If you give different answers to the question, you will both be let go, but if you give the same response you will both be punished. You have absolutely no idea what the question will be.

You would likely not be daunted by this ordeal. After all, there is a simple way to avoid the punishment. You just have to agree to give different answers to the question. But to carry out this scheme, you must do more than just agree to give different answers, you must agree, while you can still communicate, exactly which answer each one will give. If your friend merely suggests that you give different answers but then leaves before saying how he will answer, then you are no better off than before. Without knowing how he will answer, you have no means to arrange your answer to be different, no matter how much you want it to be. Unless you somehow later acquire information about how your friend has answered, your strategy will be useless.

Similarly, if neither electron has a definite disposition to be deflected either up or down by the Stern-Gerlach apparatus when they separate, then it is hard to see how they can be assured of being deflected in opposite ways without some physical mechanism that makes one sensitive to what the other does. And since the electrons can be carried arbitrarily far apart, such a mechanism would have to work at arbitrary distances. Einstein never accepted the physical reality of such a mechanism, and he concluded that the electrons had to have their dispositions all along.

This conclusion in itself might seem rather obvious and mild. But everything we have said about pairs of electrons in the singlet state and $z$-oriented Stern-Gerlach magnets holds as well for the electrons and $x$-oriented magnets, or $y$-oriented magnets, or magnets oriented in any spatial direction. So if we conclude that, to avoid the spooky action-at-a-distance, each electron must have a definite disposition about how it will behave if confronted with a $z$-oriented magnet, then it must equally have a definite disposition with respect to $x$-oriented magnets, $y$-oriented magnets, and so on. But we have already seen that we can't prepare a beam of electrons so that we can both predict with certainty how each electron will react to a $z$-oriented magnet and how it will react to an $x$-oriented magnet—that impossibility is an example of the Heisenberg uncertainty relation. Nonetheless, if we are to avoid Einstein's spooky action-at-a-distance, each electron in a singlet state must have a definite propensity to react a particular way to a $z$-oriented magnet and to an $x$-oriented magnet.

There is no contradiction in saying that on one hand, it is impossible to prepare a beam of electrons so that all will be deflected up if confronted with a *z*-oriented magnet and all will be deflected down if confronted with an *x*-oriented magnet, while on the other hand insisting that individual electrons have both these propensities. But if such individual electrons exist, standard quantum theory does not have the resources to represent their physical state. That was the issue as the EPR paper presented it: Is the quantum-mechanical description of a system complete (i.e., does it somehow represent all physical characteristics of the system)? Having rejected the spooky action-at-a-distance, EPR conclude that some physical characteristic of each electron must determine how it would behave in all these different experimental conditions, and therefore the quantum-mechanical description of the individual system is not complete. But as far as logic goes, one could reject their conclusion by embracing the notion of action-at-a-distance.

Einstein did not imagine that his rejection of action-at-a-distance could be subject to experimental test. In 1964, John Bell proved him wrong.

## Experiment 8: GHZ/Tests of Bell's Inequality

We have arrived at the strangest and most counterintuitive phenomena predicted by quantum theory and confirmed in the lab. We will mention two related examples of the general phenomenon, one conceptually simpler but experimentally harder, the other experimentally easier to confirm but slightly more complicated to analyze.

The conceptually simpler example was discovered in 1989 by Daniel Greenberger, Michael Horne and Anton Zeilinger, inspired by reflection on Bell's work. The experimental situation they envisage bears obvious similarities to Bohm's spin version of the EPR example, except three particles are involved rather than two. This triple of particles is created in a particular quantum-mechanical state and allowed to separate to arbitrary distances

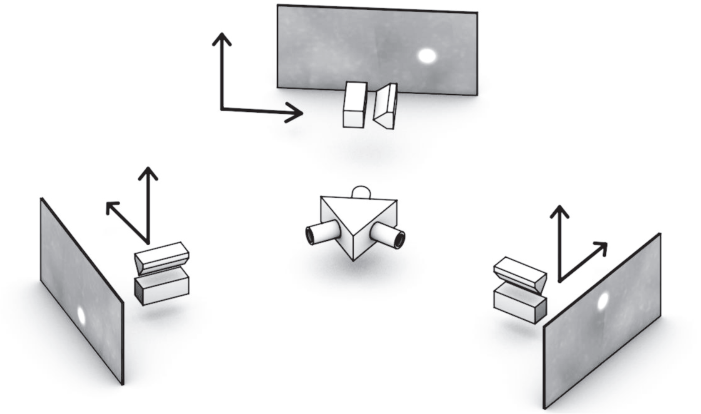

Figure 12

apart, where each will either be subjected to an *x*-oriented magnet or to a *z*-oriented magnet (Figure 12). We can imagine the choice between the two experimental arrangements for each particle being made at random in whatever way one likes: by whim, by a random number generator, by flipping a coin, and so forth. The predictions are independent of how this choice happens to be made. Figure 12 depicts two *z*-oriented and one *x*-oriented magnets.

If we denote the particles by 1, 2, and 3, then the possible local experimental conditions can be labeled $X_1$, $Z_1$, $X_2$, $Z_2$, $X_3$ and $Z_3$. The global experimental situation in a particular run of the experiment will specify the situation for each of the three magnets, so there will be eight possible global experimental configurations: $X_1X_2X_3$, $X_1X_2Z_3$, $X_1Z_2X_3$, $X_1Z_2Z_3$, $Z_1X_2X_3$, $Z_1X_2Z_3$, $Z_1Z_2X_3$, and $Z_1Z_2Z_3$. If we decide which way to set each apparatus by the flip of a fair coin, then we would expect each of these global conditions to obtain about once in every eight runs of the experiment.

Of these eight possible global configurations, currently we are only interested in four: $X_1X_2X_3$, $X_1Z_2Z_3$, $Z_1X_2Z_3$, and $Z_1Z_2X_3$. After many runs of the experiment, we would notice the following

(continued...)

# Index