

entropy nH . Let N be the minimum number of bits needed to represent a message block. This will be the smallest integer that is at least as big as nH , and so

$$nH \leq N < nH + 1. \quad (1.5)$$

Calculating the number of bits required on a “per message” basis, we are using $K = N/n$ bits per message, and

$$H \leq K < H + \frac{1}{n}. \quad (1.6)$$

If we consider very large message blocks, $n \gg 1$ and so $1/n$ is very small. The two ends of the inequality chain squeeze together, and for large blocks we will use almost exactly H bits per message to represent the information. Therefore, if we encode our messages “wholesale,” the entropy H precisely measures the number of bits per message that we need.

Exercise 1.6 Consider a type of message that has three possible values (like the message of the colonial spies in Boston). Calculate the minimum number of bits required to encode blocks of 2, 3, 5, 10, or 100 such messages. In each case, also calculate the number of bits used per message.

Things become more complicated in the presence of noise. *Noise* is a general term for any process that prevents a signal from being transferred and read unambiguously. For example, imagine that there had been fog on Boston Harbor on that April night in 1775. In a heavy fog, the church steeple might not have been visible at all from Charlestown, and no information would have been conveyed. In a lighter mist, the observers might have been able to see that there were lamps in the steeple, but not been able to count them. They would then have known that the British troops were on the move, but not which way they were going. A part of the information would have been transmitted successfully, but not all.

It is possible to formalize this notion of partial information. Before any communication takes place, there are M possible messages and the entropy is $H = \log M$. Afterward, we have reduced the number of possible messages from M to M' , but because of noise $M' > 1$. The amount of information conveyed in this process is defined to be

$$H - H' = \log \frac{M}{M'}. \quad (1.7)$$

Exercise 1.7 A friend is thinking of a number between 1 and 20 (inclusive). She tells you that the number is prime. How much information has she given you?

The concept of information is fundamental in scientific fields ranging from molecular biology to economics, not to mention computer science, statistics, and various branches of engineering. It is also, as we will see, an important unifying idea in physics.

1.2 Wave-particle duality

Since the 17th Century, there have been two basic theories about the physical nature of light. Isaac Newton believed that light is composed of huge numbers of particle-like “corpuscles.” Christiaan Huygens favored the idea that light is a wave phenomenon, a moving periodic

disturbance analogous to sound. Both theories explain the obvious facts about light, though in different ways. For example, we observe that two beams of light can pass through one another without affecting each other. In the Newtonian corpuscle theory, this simply means that the light particles do not interact with each other. In the Huygensian wave theory, it implies that light waves obey the *principle of superposition*: the total light wave is simply the sum of the waves of the two individual beams.

To take another example, we observe that the shadows of solid objects have sharp edges. This is easily explained by the Newtonian theory, since the light particles move in straight lines through empty space. On the other hand, this observation seems at first to be a fatal blow to the wave theory, because waves moving past an obstacle should spread out in the space beyond. However, if the wavelength of light were very short, then this spreading might be too small to notice. For over a hundred years, the known experimental facts about light were not sufficient to settle whether light was a particle phenomenon or a wave phenomenon, and both theories had many adherents.

Then, in 1801, Thomas Young performed a crucial experiment in which Huygens's wave theory was decisively vindicated. This was the famous two-slit experiment.

Suppose that a beam of monochromatic light shines on a barrier with a single narrow opening, or "slit." The light that passes through the slit falls on a screen some distance away. We observe that the light makes a small smudge on the screen. (For thin slits, this smudge of light actually gets wider when the slit is made narrower, and on either side of the main smudge there are several much dimmer smudges. These facts are already difficult to explain without the wave theory, but we will skip this point for now.)

Light passing through another slit elsewhere in the barrier will make a similar smudge centered on a different point. But suppose two nearby slits are both open at once. If we imagine that light is simply a stream of non-interacting Newtonian corpuscles, we would expect to see a somewhat broader and brighter smudge of light, the result of the two corpuscle-showers from the individual slits.

But what happens in fact (as Young observed) is that the region of overlap of the two smudges shows a pattern of light and dark bands called *interference fringes*, see Fig. 1.1.

This is really strange. Consider a point on the screen in the middle of one of the dark fringes. When either one of the slits is open, some light does fall on this point. But when both slits are open, the spot is dark. In other words, we can *decrease* the intensity of light at some points by *increasing* the amount of light that passes through the barrier.




Fig. 1.1

The light patterns from two single slits combine to form a pattern of interference fringes. (For clarity on the printed page, the negative of the pattern is shown; more ink means higher intensity.)

The situation is no less peculiar for the bright fringes. Take a point in the middle of one of these. When either slit is opened, the intensity of light at the point has some value I . But with both slits open, instead of an intensity $2I$ (as we might have expected), we see an intensity of $4I$! The *average* of the intensity over the light and dark fringes is indeed $2I$, but the pattern of light on the screen is less uniform than a particle theory of light would suggest.

Young realized that this curious behavior could easily be explained by the wave theory of light. Waves emerge from each of the two slits, and the combined wave at the screen is just the sum of the two disturbances. Denote by $\phi(\vec{r}, t)$ the quantity that describes the wave in space and time. In sound waves, for example, the “wave function” ϕ describes variations in air pressure. The two slits individually produce waves ϕ_1 and ϕ_2 , and by the principle of superposition the two slits together produce a combined wave $\phi = \phi_1 + \phi_2$.

Two further points complete the picture. First we note that ϕ can take on either positive or negative values. By analogy to surface waves on water, the places where ϕ is greatest are called the wave “crests,” while the places where ϕ is least (most negative) are called the wave “troughs.” Second, the observed intensity of the wave at any place is related to the square of the magnitude of the wave function there: $I \propto |\phi|^2$.

At some points on the screen, the two partial waves ϕ_1 and ϕ_2 are “out of phase,” so that a crest of ϕ_1 is coincident with a trough of ϕ_2 and vice versa. At these points, the waves cancel each other out, and $|\phi|^2$ is small. This phenomenon is called *destructive interference* and is responsible for the dark fringes.

At certain other points on the screen, the two partial waves ϕ_1 and ϕ_2 are “in phase,” by which we mean that their crests and troughs arrive synchronously. When ϕ_1 is positive, so is ϕ_2 , and so on. The partial waves reinforce each other, and $|\phi|^2$ is large. This phenomenon, *constructive interference*, is responsible for the bright fringes.

At intermediate points, ϕ_1 and ϕ_2 neither exactly reinforce one another nor exactly cancel, so the resulting intensity has an intermediate value.

Exercise 1.8 In the two slit experiment, in a particular region of the screen the light from a single slit has an intensity I , but when two slits are open, the intensity ranges over the interference fringes from 0 to $4I$. Explain this in terms of ϕ_1 and ϕ_2 .

Young was able to use two-slit interference to determine the wavelength λ of light, which does turn out to be quite small. (For green light, λ is only 500 nm.) Later in the 19th Century, James Clerk Maxwell put the wave theory of light on a firm foundation by showing that light is a travelling disturbance of electric and magnetic fields – an *electromagnetic wave*.

But the wave theory of light was not the last word. In the first years of the 20th Century, Max Planck and Albert Einstein realized that the interactions of light with matter can only be explained by assuming that the energy of light is carried by vast numbers of discrete light *quanta* later called *photons*. These photons are like particles in that each has a specific discrete energy E and momentum p , related to the wave properties of frequency f and wavelength λ :

$$\begin{aligned} E &= hf, \\ p &= \frac{h}{\lambda}, \end{aligned} \tag{1.8}$$

where $h = 6.626 \times 10^{-34}$ J s, called *Planck's constant*. When matter absorbs or emits light, it does so by absorbing or creating a whole number of photons.

Einstein used this idea to explain the photoelectric effect. In this phenomenon, light falling on a metal in a vacuum can cause electrons to be ejected from the surface. If the light intensity is increased, the number of ejected electrons increases, but the kinetic energy of each photoelectron remains the same. In a simple wave theory, this is hard to understand. Why should a more intense light, with stronger electric and magnetic fields, not produce more energetic photoelectrons? Einstein reasoned that each ejected electron gets its energy from the absorption of one photon. A brighter light has more photons, but each photon still has the same energy as before.

Exercise 1.9 The “work function” W of a metal is the amount of energy that must be added to an electron to free it from the surface. Write down an expression for the kinetic energy K of a photoelectron in terms of W and the incident light frequency f . Also find an expression for the minimum frequency f_0 required for the photoelectric effect to take place. (This will depend on W , and so may be different for different metals.)

This “quantum theory” of light poses some perplexities. In view of Young’s two-slit interference experiment, there can be no question of abandoning the wave theory entirely. Photons cannot be Newtonian corpuscles. Nevertheless, the fact that light propagates through space as a continuous wave (as seen in the two-slit experiment) does not prevent light from interacting with matter as a collection of discrete particles (as in the photoelectric effect). Furthermore, this bizarre situation is not limited to light. In 1924 Louis De Broglie discovered that the particles of matter – electrons and so forth – also have wave properties, with particle and wave quantities related by Eq. 1.8. It is possible to do a two-slit experiment with electrons and observe interference effects. The general principle that everything in nature has both wave and particle properties is sometimes called *wave – particle duality*.

The effort to put quantum ideas into a solid, consistent mathematical theory led to the development of *quantum mechanics* by Werner Heisenberg, Erwin Schrödinger, and Paul Dirac. Quantum mechanics has proved to be a superbly successful theory of phenomena ranging from elementary particles to solid state physics. It is also a very peculiar theory that challenges our intuitions on many levels. Quantum mechanics involves far-reaching alterations in our ideas about mechanics, probability theory, and even (as we shall see) the concept of information.

To illustrate this in a small way, let us re-examine Young’s two-slit experiment with quantum eyes. First, we must understand that the intensity of light is a statistical phenomenon. When we say that light is more intense at one point than it is at another, we simply mean that more photons can be found there. But what can this mean when the number is very small? What can it mean if there is only one photon present?

In the single-photon case, the intensity of the wave at any point is proportional to the *probability* of finding the photon at that point. In general, quantum mechanics predicts only the probability of an event, not whether or not that event will definitely occur. So it is with photons. The behavior of any particular photon cannot be predicted exactly, but the




Fig. 1.2 Photons fall randomly on a screen according to a probability distribution given by two-slit interference. Each image shows four times as many photons as the one before. After many photons, a smooth intensity pattern emerges statistically.

statistical behavior of a great many photons gives rise to a smooth intensity pattern. See Fig. 1.2 for an illustration of this.

In the single-photon case, therefore, the wave ϕ is actually a *probability amplitude*, a curious mathematical creature that is not itself a probability, but from which a probability may be calculated. Roughly speaking, the probability³ P of finding a photon at a given point is just $P = |\phi|^2$. Probability is the square of the magnitude of a probability amplitude.

The probability amplitude wave ϕ obeys the principle of superposition. In the two-slit experiment, consider a particular point X on the screen. With only slit #1 open, the probability amplitude that the photon lands at X is ϕ_1 , so that the probability of finding the photon there is $P_1 = |\phi_1|^2$. Opening only slit #2 yields an amplitude ϕ_2 , which gives rise to a probability P_2 of finding the photon at X . But with both slits open, we have a combined probability amplitude $\phi = \phi_1 + \phi_2$, yielding a probability

$$P = |\phi|^2 = |\phi_1 + \phi_2|^2, \quad (1.9)$$

for the photon to wind up at X . The two probability amplitudes may reinforce one another or cancel each other out, enhancing or suppressing the probability that the photon lands at X .

If the photon can pass through only one slit, the probability of reaching X is P_1 . If it can pass only through the other, it is P_2 . In ordinary probability theory, if there are two possible mutually exclusive ways that an event can happen, then the combined probability is $P = P_1 + P_2$. For example, if we flip two coins, the probability that they land with the same face upward is

$$P(\text{same face}) = P(\text{both heads}) + P(\text{both tails}). \quad (1.10)$$

³ In the two-slit experiment, where the photon can be found in a continuous range of positions, P is actually a probability *density* rather than a probability. This technical detail, and a great many others, will be worked out carefully in later chapters!

But quantum probabilities are not ordinary probabilities! In the two slit experiment, the combined likelihood may be either less than or greater than the sum $P_1 + P_2$, depending on the relative phase of the two amplitudes ϕ_1 and ϕ_2 . In other words, quantum probabilities can exhibit destructive and constructive interference effects.

Suppose at a point X on the screen the probabilities P_1 and P_2 both equal p . This means that the probability amplitudes at this point satisfy

$$|\phi_1| = |\phi_2| = \sqrt{p}. \quad (1.11)$$

If the two amplitudes constructively interfere at X , then the two amplitudes are “in phase” there: $\phi_1 = \phi_2$, and so

$$P = |\phi|^2 = |2\phi_1|^2 = 4p. \quad (1.12)$$

If the two amplitudes destructively interfere at X , then $\phi_1 = -\phi_2$ (the amplitudes are “out of phase”). Then $\phi = 0$ and so $P = 0$. We can see that the probability P for finding a photon in this region of the screen will vary over the interference fringes between 0 and $4p$.

Exercise 1.10 Consider a point X on the screen at which $P_1 = p$ and $P_2 = 2p$. That is, with only slit #1 open, the photon has a probability p of reaching X , but with only slit #2 open this probability is twice as great. Now open both slits. What are the largest and smallest possible values for P at X due to interference effects?

When analyzing the behavior of a photon in the two-slit experiment, we find that $P = |\phi_1 + \phi_2|^2$. Yet the conventional probability law $P = P_1 + P_2$ does apply to the two-coin example. So we are faced with an apparent inconsistency. Sometimes we must add probabilities, and sometimes we must add probability amplitudes. How do we know which of these rules will apply in a given situation?

The difference cannot be mere size. Quantum interference effects have been observed in surprisingly large systems, including molecules more than a million times more massive than electrons (see [Problem 1.4](#)). Conversely, we can often apply ordinary probability rules to microscopic systems. The essential difference between the two situations must lie elsewhere.

Notice that, in the two-coin example, we can check to see which of the two contributing alternatives actually occurred. That is, we can examine the coins and tell whether they are both heads or both tails. But in the two-slit experiment, this is not possible. If the single photon arrives in one of the bright interference fringes, it could have passed through either of the slits. Even a very close examination of the apparatus afterward would not tell us which possible alternative occurred.

Suppose we were to modify the two-slit experiment so that we could tell which slit the photon passed through. We can for instance imagine a very sensitive photon detector placed beside one of the slits, which is able to register the passage of a photon without destroying it. This detector need not be a large device: a single atom would be enough in principle, if the state of that atom were sufficiently affected by the passing light quantum. With such a detector in place, we could perform the two-slit interference experiment and then afterwards determine which path the photon took, simply by checking whether or not a photon had been detected.

But, as Niels Bohr pointed out, this new experiment is *not* the same as the original two-slit experiment. If we analyze the proposed modification carefully, we will find that the presence of the detector modifies the behavior of the light. The consistent phase relationship between the partial waves from the two slits will be destroyed, and so no consistent interference effects will be observable. The pattern of light intensity (photon probability) on the screen will show no bright and dark interference fringes. In fact, the probability P of a photon arriving at a point X will be exactly the sum $P_1 + P_2$ for this experiment.

Exercise 1.11 Suppose that a particle detector is placed beside slit #2 in the two-slit experiment. As a simplified model, imagine that the effect of the detector on the quantum amplitude is to randomly multiply the partial wave ϕ_2 by $+1$ or -1 . Show that, on average, the ordinary probability law holds – that is, that the average of $|\phi_1 + \phi_2|^2$ and $|\phi_1 - \phi_2|^2$ is exactly $P_1 + P_2$. (This is true whether the amplitudes are real or complex quantities.)

Bohr said that the interference experiment and the “which slit” experiment are *complementary* measurement procedures. We can do either of them, but choosing to perform one logically excludes performing the other on the same photon. We can *either* arrange the apparatus so that interference effects are present, *or* we can arrange it so that we find out which slit the photon passed, *but not both*.

The essential difference between the two-coin experiment (sum the probabilities) and the two-slit experiment (sum the amplitudes) is *information*. In each situation, two alternatives contribute to a final result. For the coins, there is no obstacle to obtaining information about which of the two possible alternatives (heads or tails) is realized. In that case, the total probability is given by $P = P_1 + P_2$. But for a photon in a two-slit interference experiment, such information is not available. Indeed, *it does not exist*, because any actual arrangement in which the photon’s path is registered will show no interference effects at all, even if the information is never read by a human experimenter. The quantum rule for adding probability amplitudes applies when the system is *informationally isolated* and produces no physical record of any sort anywhere in the Universe about which possible intermediate alternative is realized.

Exercise 1.12 Explain the following slogan, which might be suitable for printing on a T-shirt: *Quantum mechanics is what happens when nobody is looking*.

The idea that a photon might pass through the slits and leave *no trace at all* of its precise route is slightly disturbing and does not accord with “classical” intuitions based on Newtonian mechanics. Imagine that a Newtonian particle can travel by one of two possible paths. This particle is continually interacting with all of the other particles in the Universe. The position of the planet Saturn, say, will be minutely affected by the gravitation of the particle, which will in turn depend upon the particle’s position. Therefore, by an immensely precise determination of Saturn’s motion, we should (in principle) be able to tell which path the particle followed. In classical mechanics, no system can really be informationally isolated.

In a slightly more realistic example, the path of the photon through the slits should produce a slight lateral recoil in the barrier, and a careful determination of this recoil should in principle allow us to figure out which slit was passed. Einstein proposed just

such a thought-experiment to Bohr in the course of a years-long debate about the internal consistency of quantum theory. Bohr responded that quantum mechanics must apply to the barrier as well. The two possible final states of the barrier, which we wish to use to distinguish which slit the photon went through, do have slightly different quantum descriptions. Nevertheless, the two states are not reliably distinguishable by any possible measurement, and so cannot be counted as distinct physical situations.⁴ So it remains true that no physical record exists of the photon's choice of slit, and the quantum probability law applies.

The concepts of information and distinguishability are at the heart of the theory of quantum mechanics. In the chapters that follow, we will develop that theory into a sophisticated mathematical structure and then apply it to many physical situations. Ideas about probability, measurement, and information will be our constant guides. Such guides will not make quantum mechanics seem less strange to our naive intuition, but they will help us begin to build a new quantum intuition, one that more nearly conforms to the strange and marvelous ways of nature.

Problems

Problem 1.1 We said that our definition of H applies when the possible messages are equally likely. Now consider a binary message in which 0 has probability $1/3$ and 1 has probability $2/3$. What value of H should we assign when the probabilities are not equal?

We determine this by “dividing” the message 1 into two messages, 1a and 1b, which are equally likely. Then the overall message has three equally likely possibilities (0,1a,1b). This message is composed of the original (0,1) message, followed (if the first message is 1) by the (1a,1b) message.

Next we *postulate* that

$$\begin{aligned} \left[\begin{array}{c} \text{entropy} \\ \text{of (0,1a,1b)} \\ \text{message} \end{array} \right] &= \left[\begin{array}{c} \text{entropy} \\ \text{of (0,1)} \\ \text{message} \end{array} \right] \\ &+ \left(\begin{array}{c} \text{probability} \\ \text{of message 1} \end{array} \right) \times \left[\begin{array}{c} \text{entropy} \\ \text{of (1a,1b)} \\ \text{message} \end{array} \right]. \end{aligned}$$

(Think about why this postulate might make sense.) This becomes

$$\log 3 = H + \frac{2}{3} \log 2 \quad \text{and thus} \quad H = \log 3 - \frac{2}{3} \log 2 \approx 0.918.$$

(a) Explain intuitively why H should be less than 1.0 in this situation.

⁴ Bohr also considered the case where a barrier of very low mass is given a sufficient “kick” that the photon's slit can be determined. But in this case, the quantum indeterminacy in the barrier's own position is enough to “wipe out” any interference effects! (We analyze a related example in [Section 10.4](#).) The Bohr–Einstein debate, with Einstein challenging and Bohr defending the principles of quantum theory and complementarity, played a vital role in clarifying the conceptual content of the quantum theory.

- (b) Calculate H if message 0 has probability $1/6$ and message 1 has probability $5/6$.
- (c) Generalize this idea to the following situation. Message 0 has a probability $p = k/n$ and message 1 has probability $q = k'/n$, where k , k' , and n are positive integers with $k + k' = n$. Find an expression for H in this case that only involves p and q .

Problem 1.2 Five cards are dealt face-down from a 52-card deck.

- (a) How many possible sets of five cards are there? How much information do we lack about the cards?
- (b) The first three are turned over and revealed. Knowing these, how many possibilities remain?
- (c) How much information was conveyed when the three cards were revealed? Is this $3/5$ of the total? Why or why not?
- (d) Repeat parts (a)–(c) if the five cards are dealt from five independent decks.

Problem 1.3 In his short story “The Library of Babel,” Jorge Luis Borges imagines a seemingly infinite library containing books of random text. The language of the library has twenty-five characters, and

... each book is of four hundred and ten pages; each page, of forty lines, each line, of some eighty letters which are black in color.

Calculate the entropy of one of the books in Borges’ library.

Problem 1.4 In 1999, a research group at the University of Vienna was able to observe quantum interference in a beam of C_{60} molecules. C_{60} is called *buckminsterfullerene*, and the soccerball-shaped C_{60} molecules are sometimes called *buckyballs*. A buckyball molecule has a mass of about 1.2×10^{-24} kg.

- (a) The buckyball wavelength in the experiment was about 3 pm. How fast were the molecules moving?
- (b) What would be the wavelength of an electron moving at the same speed?

Problem 1.5 The kinetic energy K of a particle is related to its momentum p by $K = p^2/2\mu$, where μ is the particle’s mass. In a gas at absolute temperature T , the molecules have a typical kinetic energy of $3k_B T/2$. Derive an expression for the *thermal de Broglie wavelength*, a typical value for the de Broglie wavelength λ of a molecule in a gas. For helium atoms ($\mu = 6.7 \times 10^{-27}$ kg), calculate the thermal de Broglie wavelength at room temperature ($T = 300$ K) and at the boiling point of helium ($T = 4$ K).

Quantum effects become most significant in matter when the thermal de Broglie wavelength of the particles is greater than their separation. At atmospheric pressure, gas molecules are about 1–2 nm apart; in a condensed phase (liquid, solid) they are about ten times closer. How do these compare with the thermal de Broglie wavelengths you calculated for helium?

Problem 1.6 A single photon passes through a barrier with four slits and strikes a screen some distance away. Consider a point X on the screen. The probability amplitudes for reaching X via the four slits are ϕ_1 , ϕ_2 , ϕ_3 , and ϕ_4 .

- (a) What is the net probability P that the photon is found at X if no measurement is made of which slit the photon passed through?
- (b) A detector is placed by slit #4, which can register whether or not the photon passes that slit (but does not absorb the photon or deflect it). What is P in this case?
- (c) The detector is now moved to a point between slits #3 and #4 and registers whether or not the photon passes through one of these slits. However, the detector does *not* record which of these two slits the photon passes. What is P in this case?

2.1 The photon in the interferometer

This chapter introduces many of the ideas of quantum theory by exploring three specific “case studies” of quantum systems. Each is an example of a *qubit*, a generic name for the simplest type of quantum system. The concepts we develop will be incorporated into a rigorous mathematical framework in the next chapter. Our business here is to provide some intuition about why that mathematical framework is reasonable and appropriate for dealing with the quantum facts of life.

Interferometers

In [Section 1.2](#) we discussed the two-slit interference experiment with a single photon. In that experiment, the partial waves of probability amplitude were spread throughout the entire region of space beyond the two slits. It is much easier to analyze the situation in an *interferometer*, an optical apparatus in which the light is restricted to a finite number of discrete *beams*. The beams may be guided from one point to another, split apart or recombined as needed, and when two beams are recombined into one, the result may show interference effects. At the end of the interferometer, one or more sensors can measure the intensity of various beams. (A beam is just a possible path for the light, so there is nothing paradoxical in talking about a beam of zero intensity.) [Figure 2.1](#) shows the layout of a *Mach–Zehnder interferometer*, which is an example of this kind of apparatus.

What happens when we do an interferometer experiment with a single photon? We will consider this question for interferometers that contain only *linear* optical devices, which do not themselves create or absorb photons.¹ At the end of our interferometer, our light sensors are *photon detectors*, which can register the presence or absence of a single photon. Thus, in our calculations we will be interested in the probabilities that the various detectors will “click,” recording the presence of the photon in the corresponding beam.

We learned in our discussion of the two-slit experiment in [Section 1.2](#) that the probability of finding the photon at a particular location is the square of the magnitude of a probability

¹ These devices are also *unchanged* by the passage of a photon. For instance, we assume it is impossible to determine whether or not a photon has reflected from a given mirror, simply by examining the mirror afterward. The photon therefore remains *informationally isolated* during its passage through the interferometer. As we will see in [Section 10.4](#), this is an entirely reasonable assumption for actual interferometer experiments.

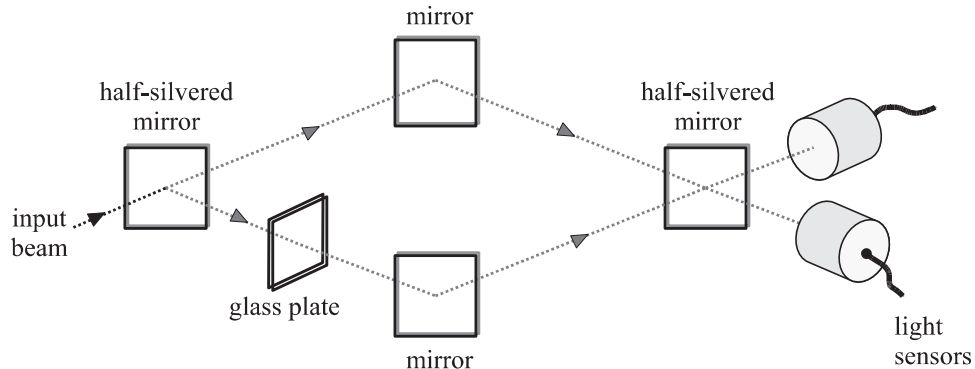


Fig. 2.1 Layout of a Mach-Zehnder interferometer. Light in the input beam is divided into two beams, which are later recombined. Light sensors measure the intensities of the two output beams.

amplitude. Each beam in our single-photon interferometer experiment will have an amplitude α , and the probability P that a detector would find the photon there (if we were to introduce such a detector) is just

$$P = |\alpha|^2. \quad (2.1)$$

Suppose at some stage of our interferometer we know for sure that the photon must be in one of two beams, which have amplitudes α and β respectively. Then it follows that $|\alpha|^2 + |\beta|^2 = 1$.

Complex amplitudes

One important kind of device that we can introduce into a beam is called a *phase shifter*. This could simply be a glass plate through which the beam travels. A phase shifter does not alter the probability that the photon is found in the beam, so the magnitude $|\alpha|$ is not changed. However, the *phase* of α can be altered. By introducing a particular thickness δ of glass, we can change the amplitude from α to $-\alpha$. (The exact value of δ depends on the index of refraction of the glass and the wavelength of the light.) This change in phase is highly significant, for it can turn constructive interference into destructive interference at a later stage of the interferometer.

If we have two such plates, or a single plate with thickness 2δ , the amplitude will become $-(-\alpha) = \alpha$, and the original amplitude is restored. But suppose we have a plate of thickness $\delta/2$? This plate would produce a change the amplitude α such that (1) the magnitude $|\alpha|$ is still the same, and (2) if the change were performed twice, the phase would be multiplied by -1 .

Glass plates can be made in a continuous range of thicknesses, producing a continuous range of phase shifts. For this to be possible, *the beam phases α must be complex quantities, with both real and imaginary parts*. A plate with thickness $\delta/2$ may multiply the amplitude by a factor of $i = \sqrt{-1}$. This does not change the magnitude of the complex phase α , since $|\alpha| = |i\alpha|$. Two such plates (or a single plate of thickness δ) multiply the phase by $i^2 = -1$, as required.




Fig. 2.2 Two important interferometer components. The photon detector D will register the presence of a photon in the beam with probability $P = |\alpha|^2$, where α is the probability amplitude. A phase shift of ϕ changes the amplitude from α to $e^{i\phi}\alpha$.

In general, a glass plate of some thickness will multiply the amplitude of the beam by $e^{i\phi}$, where ϕ (the *phase shift*) is proportional to the thickness of the glass. Changing α to $-\alpha$ could be accomplished by phase shifters with $\phi = \pi, 3\pi, 5\pi$, and so on. A phase shift of ϕ does not change the probability that the photon is found in the beam, since for any α , see Fig. 2.2,²

$$|e^{i\phi}\alpha|^2 = |\alpha|^2. \quad (2.2)$$

The fact that quantum probability amplitudes are complex quantities is one of the oddest facts about quantum mechanics. Mathematicians introduced complex numbers in the 16th Century to help solve certain algebraic problems. Such numbers are often viewed as highly abstract entities, little connected to the physical world. The number i is, after all, said to be “imaginary.” Complex numbers are sometimes used as an algebraic shortcut in Newtonian mechanics or electromagnetism. But in quantum mechanics, complex numbers are not just a convenient trick; they are inescapable and full of significance.

Exercise 2.1 Remind yourself of the rules of complex arithmetic. If α^* denotes the complex conjugate of α , show

- $|\alpha|^2 = \alpha^*\alpha$.
- $\alpha + \alpha^* = 2\Re(\alpha)$.
- For real ϕ , $(e^{i\phi})^* = e^{-i\phi}$.

Exercise 2.2

- Suppose δ is the smallest thickness of glass that produces a phase shift of π – in other words, that multiplies the phase by -1 . What is the phase shift if the glass plate has a thickness of $\delta/5$?
- Suppose δ is the *next-to-smallest* thickness of glass that produces the same change in phase (i.e. multiplying the phase by -1). What is the smallest thickness that would do so? What phase shift would be produced by a plate of thickness $\delta/5$?

The beam amplitudes in an interferometer obey the principle of superposition. We will illustrate this with a simple example. Suppose at some stage of the interferometer, there

² Anything that changes the optical path length of the beam, including a distance of empty space, will act as a phase shifter. In our simplified treatment here, we will ignore the effect of distance and think of all phase shifters as discrete objects that can be either put into or left out of the interferometer beam.

$$\alpha \begin{pmatrix} \text{1} \longrightarrow \\ \text{0} \longrightarrow \end{pmatrix} + \beta \begin{pmatrix} \text{0} \longrightarrow \\ \text{1} \longrightarrow \end{pmatrix} = \begin{pmatrix} \alpha \longrightarrow \\ \beta \longrightarrow \end{pmatrix}$$

\uparrow
 \uparrow

situation A
situation B

Fig. 2.3 A graphical representation of Eq. 2.4, showing a superposition of situation A and situation B.

are just two beams available for the photon, which we will call the “upper” beam and the “lower” beam. Consider two possible physical situations, denoted A and B . In situation A , the photon is certainly in the upper beam. The probability amplitude for this beam is 1 and the amplitude for the lower beam is 0. (The upper beam amplitude could be anything of the form $e^{i\phi}$, but we will consider the simplest case.) In situation B , the roles are reversed: the upper amplitude is 0 and the lower is 1, and so the photon is certainly in the lower beam.

The principle of superposition means that the existence of these two situations implies the existence of many other situations in which the beam amplitudes are linear combinations of the assignments for A and B . Given complex coefficients α and β , then there is a possible physical situation which we can formally write as

$$\alpha (\text{situation } A) + \beta (\text{situation } B). \quad (2.3)$$

In this combined situation, the amplitude for the upper beam is just $\alpha \cdot 1 + \beta \cdot 0 = \alpha$, while the lower beam amplitude is $\alpha \cdot 0 + \beta \cdot 1 = \beta$. Of course, to maintain a proper assignment of probabilities, we will have to require that $|\alpha|^2 + |\beta|^2 = 1$.

This is much easier to express if we describe each situation by a column vector whose entries are the beam amplitudes. Then the first situation could be written $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and the second one $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$. The principle of superposition tells us that

$$\begin{pmatrix} \alpha \\ \beta \end{pmatrix} = \alpha \begin{pmatrix} 1 \\ 0 \end{pmatrix} + \beta \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \quad (2.4)$$

is also a possible physical situation, provided $|\alpha|^2 + |\beta|^2 = 1$, see Fig. 2.3 for an illustration. From this we note, first, that a physical situation for the photon in the interferometer can be summarized by a vector whose components are probability amplitudes. Second, the principle of superposition means that a complex linear combination of two such vectors also represents a possible physical situation, provided the amplitudes satisfy a *normalization* condition (meaning that all probabilities must add up to one).

Beamsplitters

Now we turn our attention to a key element of an interferometer, the *beamsplitter*. This is a device that takes an input beam and splits it into two beams of lower intensity. A typical beamsplitter is a half-silvered mirror. A beam incident on such a mirror will produce both a reflected beam and a transmitted beam, each having half the intensity of the original.




Fig. 2.4 At beamsplitter BS, input beams of unit amplitude produce output beams with amplitudes w , x , y , and z .




Fig. 2.5 The general situation for the beamsplitter BS. Input amplitudes α and β are transformed into output amplitudes α' and β' , each of which is a linear combination of the input amplitudes.

What is the effect of a beamsplitter on the probability amplitudes when the incident beam has only a single photon? Figure 2.4 summarizes. There are two possible input beams for the beamsplitter. For an upper input beam with amplitude 1, we denote the resulting reflected and transmitted beam amplitudes by w and x respectively. A lower input beam with amplitude 1 yields output beam amplitudes y and z , as shown. If the beamsplitter is a half-silvered mirror, then the probability that the photon is reflected or transmitted at the mirror is one-half. That is,

$$|w|^2 = |x|^2 = |y|^2 = |z|^2 = \frac{1}{2}. \quad (2.5)$$

Now we can apply the principle of superposition to find how the beamsplitter works for situations in which the photon could be in either input beam. Suppose α and β are the amplitudes for the upper and lower input beam. The beamsplitter transforms these into amplitudes α' and β' for the corresponding output beams. By superposition, these are

$$\begin{aligned} \alpha' &= w\alpha + y\beta, \\ \beta' &= x\alpha + z\beta, \end{aligned} \quad (2.6)$$

as shown in Fig. 2.5. The relation between input and output amplitudes is easy to express in the amplitude-vector notation introduced above. It is

$$\begin{pmatrix} \alpha \\ \beta \end{pmatrix} \longrightarrow \begin{pmatrix} \alpha' \\ \beta' \end{pmatrix} = \begin{pmatrix} w & y \\ x & z \end{pmatrix} \begin{pmatrix} \alpha \\ \beta \end{pmatrix}. \quad (2.7)$$

This is pretty neat. We represent the photon amplitudes by column vectors $\begin{pmatrix} \alpha \\ \beta \end{pmatrix}$ and $\begin{pmatrix} \alpha' \\ \beta' \end{pmatrix}$. The beamsplitter is described by the 2×2 matrix $\begin{pmatrix} w & y \\ x & z \end{pmatrix}$. The action of the beamsplitter on the input amplitudes then corresponds to simple matrix multiplication.

Exercise 2.3 Verify that Eq. 2.7 is correct.

So far, so good. But what are the elements of the beamsplitter matrix for a particular device? For a half-silvered mirror, we know from Eq. 2.5 that the matrix elements are complex quantities with magnitude $\frac{1}{\sqrt{2}}$. The simplest possible choice would therefore be

$$w = x = y = z = \frac{1}{\sqrt{2}}.$$

What would be the properties of such a beamsplitter? Photons incident along one or the other of the two input beams yield

$$\begin{pmatrix} 1 \\ 0 \end{pmatrix} \rightarrow \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix}$$

$$\begin{pmatrix} 0 \\ 1 \end{pmatrix} \rightarrow \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix}.$$

These are perfectly reasonable amplitudes for the output beams. In either case, the photon has a probability $\left|\frac{1}{\sqrt{2}}\right|^2 = \frac{1}{2}$ of being found in each of the output beams. But suppose we consider an input that is a superposition of the two beams:

$$\begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix} \rightarrow \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \end{pmatrix}.$$

Now the photon has probability $|1|^2 = 1$ of being found in *each* output beam. This is certainly wrong! The “simplest possible” matrix elements for a beam splitter thus cannot correspond to any actual beamsplitter, because that matrix can lead to illegal probability assignments. It does not “conserve probability.”

The output probabilities are too large because constructive interference of the amplitudes takes place in both output beams. This is not possible. If constructive interference happens in some places, destructive interference must happen elsewhere.

In other words, our “simplest possible” beamsplitter matrix fails because the phases of the matrix elements cannot be as proposed. On the other hand, this matrix works fine:

$$\begin{pmatrix} w & y \\ x & z \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}. \quad (2.8)$$

Exercise 2.4 Show that, for any allowable input amplitudes $\begin{pmatrix} \alpha \\ \beta \end{pmatrix}$, a beamsplitter described by Eq. 2.8 yields output amplitudes such that $|\alpha'|^2 + |\beta'|^2 = 1$.

Equation 2.8 describes a device called a *balanced* beamsplitter. The negative sign in the lower-right (z) matrix element means that when the lower input beam is reflected,




Fig. 2.6

A Mach-Zehnder interferometer. Compare Fig. 2.1.

it undergoes a phase shift of π , but other reflected and transmitted beams have zero net phase shift.

This accords with classical wave optics. A real half-silvered mirror is a slab of glass with a very thin metallic coating on one side. When light is reflected at an interface, the wave picks up a π phase shift whenever the incident beam is coming from a medium of lower refractive index to one of higher index – for instance, from air to glass. Thus, the beam that is reflected on the metal coating from outside the glass gets a negative sign, but not the one that reflects from the inside.³

When we include a balanced beamsplitter in our calculations, we will have to be careful to indicate on which side the reflected beam acquires the negative sign. In diagrams, we will do this by placing a dot (\bullet) on one side of the beamsplitter. The reflected beam amplitude on the dotted side is multiplied by -1 .

Consider Fig. 2.6, a diagram of the Mach-Zehnder interferometer sketched in Fig. 2.1 above. Two balanced beamsplitters BS1 and BS2 are present, as are a pair of mirrors (both labelled M) and a pair of photon detectors designated D0 and D1. A phase shifter is present on one of the beams, which introduces a phase shift of ϕ . We send photons into the interferometer along just one of the input beams, so that the amplitude of that beam can be taken to be 1.

Exercise 2.5 Consider the Mach-Zehnder interferometer set-up in Fig. 2.6, and suppose $\phi = 0$.

- Ignoring any effects of the mirrors M, show that the probabilities P_0 and P_1 of the photon being detected by D0 and D1, respectively, are just 1 and 0. In other words, there is constructive interference for D0 and destructive interference for D1.
- Is your answer in part (a) changed if you take into account that reflection from a mirror M introduces a phase shift of π into that beam?

See also Problem 2.1.

³ For simplicity we are neglecting other phase shifts due to the thickness of the glass. However, if these are arranged to be integer multiples of 2π , or if the beamsplitter is built so that all beams undergo exactly the same phase shifts, these may be ignored.

Matrix methods

At any stage of a Mach–Zehnder interferometer, the photon may be in one of two possible beams. We have drawn our diagrams so that one beam is the “upper” beam and one is the “lower” beam. Devices such as phase shifters and beamsplitters alter the probability amplitudes of those beams in a linear way. This linearity is what permits us to describe the transformations by matrices.

The physical situation is described by a column vector of probability amplitudes:

$$\mathbf{v} = \begin{pmatrix} \alpha \\ \beta \end{pmatrix}. \quad (2.9)$$

The various elements of an interferometer apparatus are described by matrices acting on the amplitude vector \mathbf{v} . The balanced beamsplitter of Eq. 2.8 is described by:

$$\mathbf{B}_l = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}. \quad (2.10)$$

The subscript l indicates that the negative phase appears when the lower beam is reflected. This beamsplitter transforms the amplitude vector \mathbf{v} to a new vector \mathbf{v}' according to

$$\mathbf{v}' = \mathbf{B}_l \mathbf{v}. \quad (2.11)$$

A phase shifter can also be described by a matrix. Suppose the phase of the upper beam is shifted by ϕ . This can be represented by the matrix

$$\mathbf{P}_u(\phi) = \begin{pmatrix} e^{i\phi} & 0 \\ 0 & 1 \end{pmatrix}, \quad (2.12)$$

and the amplitude vector transforms by $\mathbf{v}' = \mathbf{P}_u(\phi)\mathbf{v}$. Once again, the subscript u indicates that the phase of the upper beam is shifted.

Exercise 2.6 Write down the matrices \mathbf{B}_u and $\mathbf{P}_l(\phi)$ describing a beamsplitter with the opposite orientation (negative phase for upper beam reflection) and a phase shifter on the lower beam.

The full-silvered mirrors that guide the beam around the interferometer introduce phase shifts by π into the beam, so they can be represented by matrices

$$\mathbf{M}_{u,l} = \mathbf{P}_{u,l}(\pi). \quad (2.13)$$

We finish our inventory with two very simple cases. First, we can imagine an arrangement in which the beams are simply allowed to cross one another, without any beamsplitter intervening. This just exchanges the upper and lower amplitudes, and so can be represented by the matrix

$$\mathbf{X} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}. \quad (2.14)$$

Simplest of all is a part of the interferometer in which the beams are not affected by any sort of optical element, and the amplitudes are unchanged. This is a sort of “device” as well! Its (trivial) action is represented by the identity matrix:

$$\mathbf{1} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}. \quad (2.15)$$




Fig. 2.7

Representations of various linear optical elements in an interferometer.




Fig. 2.8

The Mach-Zehnder interferometer. Compare Fig. 2.6.

Obviously, $\mathbf{1}\mathbf{v} = \mathbf{v}$ for any amplitude vector \mathbf{v} .

We can represent each of these graphically using a modification of our previous diagrams. From now on we will draw the upper and lower beam paths as parallel lines, except where they are brought together at a beamsplitter or a beam crossing. The photon is assumed to go from left to right, see Fig. 2.7.⁴

What happens when the basic optical elements are assembled into a larger experiment? In a diagram, we simply string the pieces together in sequence, as in Fig. 2.8. How can we describe this sort of interferometer arrangement mathematically? Suppose a pair of beams with amplitude vector \mathbf{v} pass through three optical elements. The first is described by a matrix \mathbf{R} , the second by \mathbf{S} , and the third by \mathbf{T} . To find the final amplitude vector \mathbf{v}' , we must first multiply \mathbf{v} by \mathbf{R} , then by \mathbf{S} , then by \mathbf{T} :

$$\mathbf{v}' = \mathbf{TSR}\mathbf{v}. \quad (2.16)$$

The effect of the entire complex apparatus is represented by a *single* 2×2 matrix, the product \mathbf{TSR} . This product is a sequence in time of successive transformations of the amplitude vector for the beams, with the *time order* from right to left: \mathbf{R} occurs first and \mathbf{T} occurs last. To put it another way, the order of the matrices in the product is the opposite of the order of the corresponding elements in our left-to-right diagrams.

Exercise 2.7 Write down a matrix product that represents the Mach-Zehnder interferometer shown in Fig. 2.8. (You may ignore the photon detectors at the end.)

⁴ Do not be worried by the fact that our beams no longer go in straight lines in our diagrams. The diagrams are merely schematics of a real optical apparatus. But as a matter of fact, we can build interferometers in which the beams are guided in curved paths by optical fibers.

As you thought about [Exercise 2.7](#), you may have noticed a difficulty. The two beams strike two different mirrors, each of which yields a phase shift of π . These reflections happen at about the same time, as suggested in [Fig. 2.8](#). In which order should we write the corresponding matrices? Fortunately, it turns out that the order of these phase shifter matrices does not matter. We will cast the relevant fact as an exercise:

Exercise 2.8 Suppose \mathbf{P} and \mathbf{P}' are the matrices for two phase shifters. Show that \mathbf{P} and \mathbf{P}' commute:

$$\mathbf{P}\mathbf{P}' = \mathbf{P}'\mathbf{P}$$

when (a) the two phase shifters are applied to the same beam, and (b) the two phase shifters are applied to different beams.

Some of the matrices commute with each other, but not all of them. For example:

Exercise 2.9 Show that

$$\mathbf{X}\mathbf{P}_u(\pi) \neq \mathbf{P}_u(\pi)\mathbf{X}.$$

Explain in words why this makes sense.

The analysis of a two-beam interferometer system has now been boiled down to matrix calculations. The translation between the physical apparatus and the mathematical expression is straightforward. The following exercise should give you some easy practice at these calculations and translations. You will find more examples in the problems at the end of the chapter.

Exercise 2.10 Verify the following matrix facts, and explain each one in words and pictures as a fact about interferometer systems. (a) $\mathbf{X}\mathbf{X} = \mathbf{1}$. (b) $\mathbf{B}_l\mathbf{B}_l = \mathbf{1}$. (c) $\mathbf{B}_l\mathbf{P}_l(\pi)\mathbf{P}_u(\pi)\mathbf{B}_l = -\mathbf{1}$. (d) $\mathbf{B}_l\mathbf{P}_l(\pi)\mathbf{B}_l = \mathbf{X}$. (e) $\mathbf{B}_l\mathbf{P}_l(\pi)\mathbf{B}_u = \mathbf{P}_u(\pi)$.

Because of the principle of superposition, any linear optical element will produce a linear transformation on the input amplitude vector \mathbf{v} , and can therefore be represented by a 2×2 matrix \mathbf{R} acting on \mathbf{v} . But we saw in our analysis of beamsplitters that not all 2×2 matrices could possibly correspond to an actual optical device. The reason was that some matrices did not preserve the normalization of the probabilities. Which matrices \mathbf{R} do preserve this normalization, and so might correspond to actual devices?

First, we need to express the normalization requirement in terms of matrices. The *Hermitian conjugate* operation is designated by the “dagger” symbol “ \dagger ”. This indicates the complex conjugate of the transpose of the matrix. Thus,

$$\mathbf{v}^\dagger = \begin{pmatrix} \alpha^* & \beta^* \end{pmatrix}. \quad (2.17)$$

Our normalization requirement for the probability amplitudes can then be written as

$$\mathbf{v}^\dagger \mathbf{v} = 1. \quad (2.18)$$

(Note that we are equating the number 1 with the 1×1 matrix whose only entry is 1. This is a harmless abuse of mathematical notation.)

Exercise 2.11 Verify that this equation is the same as $|\alpha|^2 + |\beta|^2 = 1$.

The vector $\mathbf{v}' = \mathbf{R}\mathbf{v}$ contains the output amplitudes when the input is \mathbf{v} . We are thus requiring that $(\mathbf{v}')^\dagger \mathbf{v}' = 1$ for any input vector that has $\mathbf{v}^\dagger \mathbf{v} = 1$. In other words,

$$(\mathbf{v}')^\dagger \mathbf{v}' = \mathbf{v}^\dagger \mathbf{R}^\dagger \mathbf{R} \mathbf{v} = 1. \quad (2.19)$$

(We have used the fact that, for any complex matrices, $(\mathbf{UV})^\dagger = \mathbf{V}^\dagger \mathbf{U}^\dagger$. This, or at least the corresponding fact for the matrix transpose, should be familiar.)

We can view Eq. 2.19 as a property of the matrix $\mathbf{R}^\dagger \mathbf{R}$. Let

$$\mathbf{R}^\dagger \mathbf{R} = \begin{pmatrix} q & r \\ s & t \end{pmatrix}. \quad (2.20)$$

What can we say about these matrix elements? First, consider an input amplitude vector $\mathbf{v} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$. Then

$$\mathbf{v}^\dagger \mathbf{R}^\dagger \mathbf{R} \mathbf{v} = \begin{pmatrix} 1 & 0 \end{pmatrix} \begin{pmatrix} q & r \\ s & t \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} = q. \quad (2.21)$$

So Eq. 2.19 tells us that $q = 1$.

Exercise 2.12 Verify Eq. 2.21, and then repeat the calculation with $\mathbf{v} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$ to show that $t = 1$.

The two diagonal elements of $\mathbf{R}^\dagger \mathbf{R}$ must both equal 1. What about the other two elements?

If we let $\mathbf{v} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix}$, we have

$$\mathbf{v}^\dagger \mathbf{R}^\dagger \mathbf{R} \mathbf{v} = \frac{1}{2} \begin{pmatrix} 1 & 1 \end{pmatrix} \begin{pmatrix} 1 & r \\ s & 1 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \end{pmatrix} = 1 + \frac{1}{2}(r + s). \quad (2.22)$$

Since this must equal 1, we know that $s = -r$. Finally, we recall that the amplitudes are complex numbers, so that the input $\mathbf{v} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ i \end{pmatrix}$ is possible. This yields

$$\mathbf{v}^\dagger \mathbf{R}^\dagger \mathbf{R} \mathbf{v} = \frac{1}{2} \begin{pmatrix} 1 & -i \end{pmatrix} \begin{pmatrix} 1 & r \\ -r & 1 \end{pmatrix} \begin{pmatrix} 1 \\ i \end{pmatrix} = 1 + ri. \quad (2.23)$$

From this, we conclude that $r = 0$.

Exercise 2.13 Verify Eq. 2.22 and Eq. 2.23.

Putting it all together, we have shown that, if the matrix \mathbf{R} is to preserve the normalization of probabilities, it must have the property that

$$\mathbf{R}^\dagger \mathbf{R} = \mathbf{1}. \quad (2.24)$$

Matrices with this property are called *unitary* matrices. We have arrived at an important general fact: *Any physically possible linear optical element in a two-beam interferometer is represented by a 2×2 unitary matrix.*

Exercise 2.14 Here is what we have proved: If \mathbf{R} is to preserve the normalization of probabilities for any input \mathbf{v} , then it must be unitary.

Now you prove the (much easier) converse: If \mathbf{R} is unitary, then it will preserve this normalization for any input \mathbf{v} . (Be sure that you understand the distinction between these statements!)

We can further show that any unitary 2×2 matrix \mathbf{R} may be physically realized as an interferometer set-up made out of beam splitters and phase-shifters, see [Problem 2.3](#).

Testing bombs

The components of an interferometer do not register the passage of a photon, so that the photon remains informationally isolated. This is why the beams exhibit interference. Consider, for example, the simplified Mach–Zehnder arrangement in [Fig. 2.9](#). The photon is introduced along the lower beam, so the input amplitude vector can be taken to be $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$. If nothing else is introduced into the apparatus, the matrix describing the interferometer's effect is just

$$\mathbf{B}_l \mathbf{B}_u = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}. \quad (2.25)$$

The output amplitude vector is thus

$$\mathbf{B}_l \mathbf{B}_u \mathbf{v} = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}. \quad (2.26)$$

Exercise 2.15 Check this matrix arithmetic.

Therefore, the photon will always reach the upper detector D0. The probabilities are

<i>outcome</i>	<i>P</i>
photon reaches D0	1
photon reaches D1	0.

There is constructive interference in the beam that leads to D0, and destructive interference in the beam that leads to D1.

Now suppose that we change the interferometer slightly by sticking a hand into the lower beam at the point A. For simplicity, imagine that the photon is absorbed if it hits the hand.




Fig. 2.9

Simplified Mach–Zehnder interferometer.

This produces a physical change in the hand that could in principle be detected (“Ow!”). Thus, the hand is a photon detector that measures whether or not the photon travels along the lower beam at A.

This will, of course, destroy any interference effects. If we send a photon into the apparatus, it has a 50% probability of striking the hand. If it travels along the upper beam instead, when it reaches the second beamsplitter it will be equally likely to go toward D0 and D1. In short, we have

<i>outcome</i>	<i>P</i>
photon reaches D0	1/4
photon reaches D1	1/4
photon hits hand	1/2.

Notice that, by blocking one beam with a hand, we have actually *increased* the probability that the photon is detected by D1.

This paradoxical result is the basis for a remarkable thought-experiment proposed by Avshalom Elitzur and Lev Vaidman in 1993. Imagine a factory that produces a type of bomb triggered by light. So sensitive is the trigger that the passage of a *single photon* through its mechanism will explode a bomb.

Because of manufacturing defects, however, many bombs come off the assembly line without working triggers. Photons pass through these mechanisms without being registered at all, and the bombs are duds. The factory managers want to be able tell for sure that at least some bombs are in working order. How can they do this? Of course, if they send a photon through a given bomb, and it blows up, then they can be sure that the bomb was in working order – but they have also destroyed that bomb. What the managers want is a way to identify bombs that are explosive, but are not yet exploded. Since the bomb triggers are set off even by one photon, this appears impossible.

But in fact, the interferometer arrangement in Fig. 2.9 can do the job. A bomb is placed at the point A and then one photon is sent through. If the bomb is a dud, it will not register the passage of the photon, and there will be interference effects. If the bomb is working, it will function as a photon detector on the lower path. The results are

Bomb is a dud		Bomb is working	
<i>outcome</i>	<i>P</i>	<i>outcome</i>	<i>P</i>
photon reaches D0	1	photon reaches D0	1/4
photon reaches D1	0	photon reaches D1	1/4
bomb explodes	0	bomb explodes	1/2.

Suppose an unknown bomb is placed in the apparatus and one photon is sent through. If the bomb explodes, then it was in working order, but this bomb is now lost. If the photon is detected by D0, the test is inconclusive and may be repeated.⁵ But if the photon ever arrives at D1, then the managers know that the unexploded bomb is in working order, *even though the bomb never detects the passage of the photon*.

⁵ If the photon always arrives at D0 during many trials, the factory managers may confidently conclude that the bomb is a dud.

Exercise 2.16 If you do not find the previous paragraph strange and disturbing, re-read it.

Exercise 2.17 Suppose the interferometer test is performed on a large number of bombs from the factory. When the test is inconclusive on a particular bomb, it is repeated until the bomb's status is reasonably certain. What fraction of the working bombs are certified as working but not detonated?

The Elitzur–Vaidman thought-experiment is a good example of the sometimes perplexing behavior of quantum systems. It also illustrates why information is such a key idea in quantum theory. Whether or not a working bomb actually detects a photon in a given trial, its final state (intact or exploded) provides a record of which beam the photon has traversed. That means that the photon was not informationally isolated in the apparatus, and so there can be no interference between the beams.

2.2 Spin 1/2

Having analyzed in detail the problem of a single photon in a two-beam interferometer, we are in a position to identify a few key ideas:

- At any point, the photon can be in one of two distinct beams. Linear superpositions of the beams are also possible.
- The physical situation of the photon is described by a vector \mathbf{v} of two complex probability amplitudes. If a given beam has an amplitude α , then $|\alpha|^2$ is the probability that a detector would find the photon in that beam. Normalization of probabilities means that $\mathbf{v}^\dagger \mathbf{v} = 1$.
- The effect of a linear optical device like a phase shifter or a beamsplitter is described by a matrix \mathbf{R} . The amplitude vector \mathbf{v} is changed to a new vector $\mathbf{v}' = \mathbf{R}\mathbf{v}$. The matrix \mathbf{R} must be unitary to guarantee that the final probabilities are normalized.
- Even a quantum system as simple as this can yield surprising results, as in the bomb-testing thought-experiment.

In this section, we will apply these same ideas to a quite different type of quantum system.

Particles with spin

A particle has angular momentum by virtue of its movement through space. It may also have an intrinsic angular momentum called *spin*. This term suggests an analogy to Newtonian physics, in which the angular momentum of an extended body like the Earth is due to both its translational and rotational motion. The quantum situation is a bit more subtle. Electrons, for instance, appear to be entirely point-like, without any spatial extent at all. We therefore cannot attribute the intrinsic spin of an electron to mere rotational motion.

Electrons, protons, and neutrons are all examples of *spin-1/2 particles*. Suppose we measure the z -component S_z of the spin angular momentum for one of these particles. The