

QUANTUM MECHANICS

THE THEORETICAL MINIMUM

WHAT YOU NEED *to* KNOW *to*
START DOING PHYSICS



New York Times bestselling author

LEONARD SUSSKIND
& ART FRIEDMAN

This book is the second volume of the Theoretical Minimum series. The first volume, *The Theoretical Minimum: What You Need to Know to Start Doing Physics*, covered classical mechanics, which is the core of any physics education. We will refer to it from time to time simply as Volume I. This second book explains quantum mechanics and its relationship to classical mechanics. The books in this series run parallel to Leonard Susskind's videos, available on the Web through Stanford University (see www.theoreticalminimum.com for a listing). While covering the same general topics as the videos, the books contain additional details, and topics that don't appear in the videos.

QUANTUM
MECHANICS

Also by Leonard Susskind

The Theoretical Minimum

What You Need To Know to Start Doing Physics
(with George Hrabovsky)

The Black Hole War

The Cosmic Landscape

QUANTUM
MECHANICS

The Theoretical Minimum

LEONARD SUSSKIND
and ART FRIEDMAN

BASIC BOOKS

*A Member of the Perseus Books Group
New York*

Copyright © 2014 by Leonard Susskind and Art Friedman

Published by Basic Books,
A Member of the Perseus Books Group

All rights reserved. Printed in the United States of America. No part of this book may be reproduced in any manner whatsoever without written permission except in the case of brief quotations embodied in critical articles and reviews. For information, address Basic Books, 250 West 57th Street, 15th Floor, New York, NY 10107-1307.

Books published by Basic Books are available at special discounts for bulk purchases in the United States by corporations, institutions, and other organizations. For more information, please contact the Special Markets Department at the Perseus Books Group, 2300 Chestnut Street, Suite 200, Philadelphia, PA 19103, or call (800) 810-4145, ext. 5000, or e-mail special.markets@perseusbooks.com.

Designed by Art Friedman and Leonard Susskind

Hilbert's Place drawings were created by Margaret Sloan.

A CIP catalog record for this book is available from the Library of Congress.

ISBN (hardcover): 978-0-465-03667-7

ISBN (ebook): 978-0-465-08061-8

10 9 8 7 6 5 4 3 2 1

*For our parents,
who made it all possible:
Irene and Benjamin Susskind
George and Trudy Friedman*

Contents

<i>Preface</i>	<i>xi</i>
<i>Prologue</i>	<i>xv</i>
<i>Introduction</i>	<i>xix</i>
1 Systems and Experiments	1
2 Quantum States	35
3 Principles of Quantum Mechanics	51
4 Time and Change	93
5 Uncertainty and Time Dependence	129
6 Combining Systems: Entanglement	149
7 More on Entanglement	183
8 Particles and Waves	235
9 Particle Dynamics	273
10 The Harmonic Oscillator	311
<i>Appendix</i>	<i>347</i>
<i>Index</i>	<i>353</i>

Preface

Albert Einstein, who was in many ways the father of quantum mechanics, had a notorious love-hate relation with the subject. His debates with Niels Bohr—Bohr completely accepting of quantum mechanics and Einstein deeply skeptical—are famous in the history of science. It was generally accepted by most physicists that Bohr won and Einstein lost. My own feeling, I think shared by a growing number of physicists, is that this attitude does not do justice to Einstein's views.

Both Bohr and Einstein were subtle men. Einstein tried very hard to show that quantum mechanics was inconsistent; Bohr, however, was always able to counter his arguments. But in his final attack Einstein pointed to something so deep, so counterintuitive, so troubling, and yet so exciting, that at the beginning of the twenty-first century it has returned to fascinate theoretical physicists. Bohr's only answer to Einstein's last great discovery—the discovery of *entanglement*—was to ignore it.

The phenomenon of entanglement is the essential fact of quantum mechanics, the fact that makes it so different from classical physics. It brings into question our entire un-

derstanding about what is *real* in the physical world. Our ordinary intuition about physical systems is that if we know everything about a system, that is, everything that can in principle be known, then we know everything about its parts. If we have complete knowledge of the condition of an automobile, then we know everything about its wheels, its engine, its transmission, right down to the screws that hold the upholstery in place. It would not make sense for a mechanic to say, “I know everything about your car but unfortunately I can’t tell you anything about any of its parts.”

But that’s exactly what Einstein explained to Bohr—in quantum mechanics, one can know everything about a system and nothing about its individual parts—but Bohr failed to appreciate this fact. I might add that generations of quantum textbooks blithely ignored it.

Everyone knows that quantum mechanics is strange, but I suspect very few people could tell you exactly in what way. This book is a technical course of lectures on quantum mechanics, but it is different than most courses or most textbooks. The focus is on the logical principles and the goal is not to hide the utter strangeness of quantum logic but to bring it out into the light of day.

I remind you that this book is one of several that closely follow my Internet course series, the Theoretical Minimum. My coauthor, Art Friedman, was a student in these courses. The book benefited from the fact that Art was learning the subject and was therefore very sensitive to the issues that might be confusing to the beginner. During the course of writing, we had a lot of fun, and we’ve tried to convey some

of that spirit with a bit of humor. If you don't get it, ignore it.

Leonard Susskind

When I completed my master's degree in computer science at Stanford, I could not have guessed that I'd return some years later to attend Leonard's physics lectures. My short "career" in physics ended many years earlier, with the completion of my bachelor's degree. But my interest in the subject has remained very much alive.

It appears that I have lots of company—the world seems filled with people who are genuinely, deeply interested in physics but whose lives have taken them in different directions. This book is for all of us.

Quantum mechanics can be appreciated, to some degree, on a purely qualitative level. But mathematics is what brings its beauty into sharp focus. We have tried to make this amazing body of work fully accessible to mathematically literate nonphysicists. I think we've done a fairly good job, and I hope you'll agree.

No one completes a project like this without lots of help. The people at Brockman, Inc., have made the business end of things seem easy, and the production team at Perseus Books has been top-notch. My sincere thanks go to TJ Kelleher, Rachel King, and Tisse Takagi. It was our good fortune to work with a talented copy editor, John Searcy.

I'm grateful to Leonard's (other) continuing education students for routinely raising thoughtful, provocative questions, and for many stimulating after-class conversations.

Rob Colwell, Todd Craig, Monty Frost, and John Nash offered constructive comments on the manuscript. Jeremy Branscome and Russ Bryan reviewed the entire manuscript in detail, and identified a number of problems.

I thank my family and friends for their kind support and enthusiasm. I especially thank my daughter, Hannah, for minding the store.

Besides her love, encouragement, insight, and sense of humor, my amazing wife, Margaret Sloan, contributed about a third of the diagrams and both Hilbert's Place illustrations. Thanks, Maggie.

At the start of this project, Leonard, sensing my real motivation, remarked that one of the best ways to learn physics is to write about it. True, of course, but I had no idea how true, and I'm grateful that I had a chance to find out. Thanks a million, Leonard.

Art Friedman

Prologue

Art looks over his beer and says, “Lenny, let’s play a round of the Einstein-Bohr game.”

“OK, but I’m tired of losing. This time, you be Artstein and I’ll be L-Bore. You start.”

“Fair enough. Here’s my first shot: God doesn’t play dice. Ha-ha, L-Bore, that’s one point for me.”

“Not so fast, Artstein, not so fast. You, my friend, were the first one to point out that quantum theory is inherently probabilistic. Heh heh heh, that’s a two-pointer!”

“Well, I take it back.”

“You can’t.”

“I can.”

“You can’t.”

Few people realize that Einstein, in his 1917 paper, “On the Quantum Theory of Radiation,” argues that the emission of gamma rays is governed by a statistical law.

A Professor and a Fiddler Walk into a Bar


Volume I was punctuated by short conversations between Lenny and George, fictional personas who were loosely based on two John Steinbeck characters. The setting for this volume of the Theoretical Minimum series is inspired by the stories of Damon Runyon. It's a world filled with crooks, con artists, degenerates, smooth operators, and do-gooders. Plus a few ordinary folks, just trying to get through the day. The action unfolds at a popular watering hole called Hilbert's Place.

Into this setting stroll Lenny and Art, two greenhorns from California who somehow got separated from their tour bus. Wish them luck. They will need it.

What to Bring

You don't need to be a physicist to take this journey, but you should have some basic knowledge of calculus and linear algebra. You should also know something about the material covered in *Volume I*. It's OK if your math is a bit rusty. We'll review and explain much of it as we go, especially the material on linear algebra. *Volume I* reviews the basic ideas of calculus.

Don't let our lighthearted humor fool you into thinking that we're writing for airheads. We're not. Our goal is to make a difficult subject "as simple as possible, but no simpler," and we hope to have a little fun along the way. See you at Hilbert's Place.



Introduction

Classical mechanics is intuitive; things move in predictable ways. An experienced ballplayer can take a quick look at a fly ball, and from its location and its velocity, know where to run in order to be there just in time to catch the ball. Of course a sudden unexpected gust of wind might fool him, but that's only because he didn't take into account all the variables. There is an obvious reason why classical mechanics is intuitive: humans, and animals before them, have been using it many times every day for survival. But no one ever used quantum mechanics before the twentieth century. Quantum mechanics describes things so small that they are completely beyond the range of the human senses. So it stands to reason that we did not evolve an intuition for the quantum world. The only way we can comprehend it is by rewiring our intuitions with abstract mathematics. Fortunately, for some odd reason, we did evolve the capacity for such rewiring.

Ordinarily, we learn classical mechanics first, before even attempting quantum mechanics. But quantum physics is much more fundamental than classical physics. As far as we know, quantum mechanics provides an exact description of every physical system, but some things are massive enough

that quantum mechanics can be reliably approximated by classical mechanics. That's all that classical mechanics is: an approximation. From a logical point of view, we should learn quantum mechanics first, but very few physics teachers would recommend that. Even this course of lectures—the Theoretical Minimum series—began with classical mechanics. Nevertheless, in these quantum lectures, classical mechanics will play almost no role except near the end, well after the basic principles of quantum mechanics have been explained. I think this is really the right way to do it, not just logically but pedagogically as well. That way we don't fall into the trap of thinking that quantum mechanics is basically just classical mechanics with a couple of new gimmicks thrown in. By the way, quantum mechanics is technically much easier than classical mechanics.

The simplest classical system—the basic logical unit for computer science—is the two-state system. Sometimes it's called a *bit*. It can represent anything that has only two states: a coin that can show heads or tails, a switch that is on or off, or a tiny magnet that is constrained to point either north or south. As you might expect, especially if you studied the first lecture of *Volume I*, the theory of classical two-state systems is extremely simple—boring, in fact. In this volume, we're going to begin with the quantum version of the two-state system, called a *qubit*, which is far more interesting. To understand it, we will need a whole new way of thinking—a new foundation of logic.

Lecture 1

Systems and Experiments

Lenny and Art wander into Hilbert's Place.

Art: What is this, the Twilight Zone? Or some kind of fun house? I can't get my bearings.

Lenny: Take a breath. You'll get used to it.

Art: Which way is up?

1.1 Quantum Mechanics Is Different

What is so special about quantum mechanics? Why is it so hard to understand? It would be easy to blame the “hard mathematics,” and there may be some truth in that idea. But that can't be the whole story. Lots of nonphysicists are

able to master classical mechanics and field theory, which also require hard mathematics.

Quantum mechanics deals with the behavior of objects so small that we humans are ill equipped to visualize them at all. Individual atoms are near the upper end of this scale in terms of size. Electrons are frequently used as objects of study. *Our sensory organs are simply not built to perceive the motion of an electron. The best we can do is to try to understand electrons and their motion as mathematical abstractions.*

“So what?” says the skeptic. “Classical mechanics is filled to the brim with mathematical abstractions—point masses, rigid bodies, inertial reference frames, positions, momenta, fields, waves—the list goes on and on. There’s nothing new about mathematical abstractions.” This is actually a fair point, and indeed the classical and quantum worlds have some important things in common. Quantum mechanics, however, is different in two ways:

1. *Different Abstractions.* Quantum abstractions are fundamentally different from classical ones. For example, we’ll see that the idea of a state in quantum mechanics is conceptually very different from its classical counterpart. States are represented by different mathematical objects and have a different logical structure.
2. *States and Measurements.* In the classical world, the relationship between the state of a system and the result of a measurement on that system is very straightforward. In fact, it’s trivial. The labels that describe

a state (the position and momentum of a particle, for example) are the *same* labels that characterize measurements of that state. To put it another way, one can perform an experiment to determine the state of a system. In the quantum world, this is not true. States and measurements are two different things, and the relationship between them is subtle and nonintuitive.

These ideas are crucial, and we'll come back to them again and again.

1.2 Spins and Qubits

The concept of spin is derived from particle physics. Particles have properties in addition to their location in space. For example, they may or may not have electric charge, or mass. An electron is not the same as a quark or a neutrino. But even a specific type of particle, such as an electron, is not completely specified by its location. Attached to the electron is an extra degree of freedom called its *spin*. Naively, the spin can be pictured as a little arrow that points in some direction, but that naive picture is too classical to accurately represent the real situation. The spin of an electron is about as quantum mechanical as a system can be, and any attempt to visualize it classically will badly miss the point.

We can and will abstract the idea of a spin, and forget that it is attached to an electron. The quantum spin is a system that can be studied in its own right. In fact, the quantum spin, isolated from the electron that carries it

through space, is both the simplest and the most quantum of systems.

The isolated quantum spin is an example of the general class of simple systems we call qubits—quantum bits—that play the same role in the quantum world as logical bits play in defining the state of your computer. Many systems—maybe even all systems—can be built up by combining qubits. Thus in learning about them, we are learning about a great deal more.

1.3 An Experiment

Let's make these ideas concrete, using the simplest example we can find. In the first lecture of *Volume I*, we began by discussing a very simple deterministic system: a coin that can show either heads (H) or tails (T). We can call this a two-state system, or a bit, with the two states being H and T . More formally we invent a “degree of freedom” called σ that can take on two values, namely $+1$ and -1 . The state H is replaced by

$$\sigma = +1$$

and the state T by

$$\sigma = -1.$$

Classically, that's all there is to the space of states. The system is either in state $\sigma = +1$ or $\sigma = -1$ and there is

nothing in between. In quantum mechanics, we'll think of this system as a qubit.

Volume I also discussed simple evolution laws that tell us how to update the state from instant to instant. The simplest law is just that nothing happens. In that case, if we go from one discrete instant (n) to the next ($n + 1$), the law of evolution is

$$\sigma(n + 1) = \sigma(n). \quad (1.1)$$

Let's expose a hidden assumption that we were careless about in *Volume I*. An experiment involves more than just a system to study. It also involves an apparatus \mathcal{A} to make measurements and record the results of the measurements. In the case of the two-state system, the apparatus interacts with the system (the spin) and records the value of σ . Think of the apparatus as a black box¹ with a window that displays the result of a measurement. There is also a “this end up” arrow on the apparatus. The up-arrow is important because it shows how the apparatus is oriented in space, and its direction will affect the outcomes of our measurements. We begin by pointing it along the z axis (Fig. 1.1). Initially, we have no knowledge of whether $\sigma = +1$ or $\sigma = -1$. Our purpose is to do an experiment to find out the value of σ .

Before the apparatus interacts with the spin, the window is blank (labeled with a question mark in our diagrams). After it measures σ , the window shows a $+1$ or a -1 . By

¹“Black box” means we have no knowledge of what's inside the apparatus or how it works. But rest assured, it does not contain a cat.

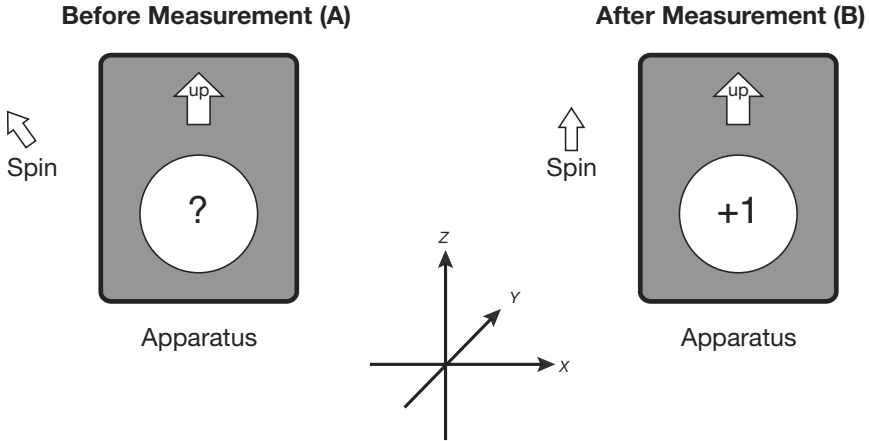


Figure 1.1: (A) Spin and cat-free apparatus before any measurement is made. (B) Spin and apparatus after one measurement has been made, resulting in $\sigma_z = +1$. The spin is now prepared in the $\sigma_z = +1$ state. If the spin is not disturbed and the apparatus keeps the same orientation, all subsequent measurements will give the same result. Coordinate axes show our convention for labeling the directions of space.

looking at the apparatus, we determine the value of σ . That whole process constitutes a very simple experiment designed to measure σ .

Now that we've measured σ , let's reset the apparatus to neutral and, without disturbing the spin, measure σ again. Assuming the simple law of Eq. 1.1, we should get the same answer as we did the first time. The result $\sigma = +1$ will be

followed by $\sigma = +1$. Likewise for $\sigma = -1$. The same will be true for any number of repetitions. This is good because it allows us to confirm the result of an experiment. We can also say this in the following way: The first interaction with the apparatus \mathcal{A} *prepares* the system in one of the two states. Subsequent experiments *confirm* that state. So far, there is no difference between classical and quantum physics.

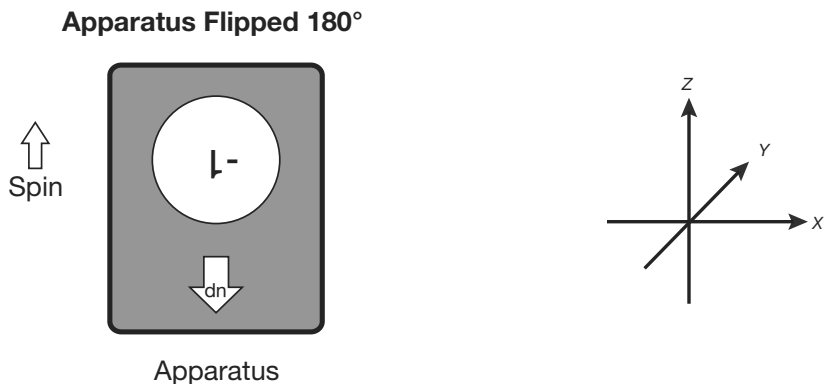


Figure 1.2: The apparatus is flipped without disturbing the previously measured spin. A new measurement results in $\sigma_z = -1$.

Now let's do something new. After preparing the spin by measuring it with \mathcal{A} , we turn the apparatus upside down and then measure σ again (Fig. 1.2). What we find is that if we originally prepared $\sigma = +1$, the upside down apparatus records $\sigma = -1$. Similarly, if we originally prepared $\sigma = -1$, the upside down apparatus records $\sigma = +1$. In other words,

turning the apparatus over interchanges $\sigma = +1$ and $\sigma = -1$. From these results, we might conclude that σ is a degree of freedom that is associated with a sense of direction in space. For example, if σ were an oriented vector of some sort, then it would be natural to expect that turning the apparatus over would reverse the reading. A simple explanation is that the apparatus measures the component of the vector along an axis embedded in the apparatus. Is this explanation correct for all configurations?

If we are convinced that the spin is a vector, we would naturally describe it by three components: σ_z , σ_x , and σ_y . When the apparatus is upright along the z axis, it is positioned to measure σ_z .

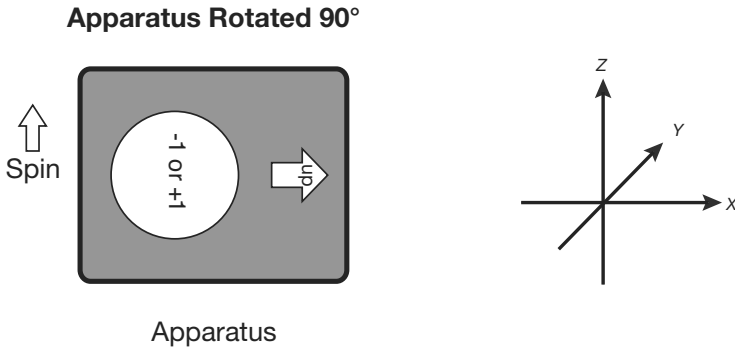


Figure 1.3: The apparatus rotated by 90°. A new measurement results in $\sigma_z = -1$ with 50 percent probability.

So far, there is still no difference between classical physics

and quantum physics. The difference only becomes apparent when we rotate the apparatus through an arbitrary angle, say $\frac{\pi}{2}$ radians (90 degrees). The apparatus begins in the upright position (with the up-arrow along the z axis). A spin is prepared with $\sigma = +1$. Next, rotate \mathcal{A} so that the up-arrow points along the x axis (Fig. 1.3), and then make a measurement of what is presumably the x component of the spin, σ_x .

If in fact σ really represents the component of a vector along the up-arrow, one would expect to get zero. Why? Initially, we confirmed that σ was directed along the z axis, suggesting that its component along x must be zero. But we get a surprise when we measure σ_x : Instead of giving $\sigma_x = 0$, the apparatus gives either $\sigma_x = +1$ or $\sigma_x = -1$. \mathcal{A} is very stubborn—no matter which way it is oriented, it refuses to give any answer other than $\sigma = \pm 1$. If the spin really is a vector, it is a very peculiar one indeed.

Nevertheless, we do find something interesting. Suppose we repeat the operation many times, each time following the same procedure, that is:

- Beginning with \mathcal{A} along the z axis, prepare $\sigma = +1$.
- Rotate the apparatus so that it is oriented along the x axis.
- Measure σ .

The repeated experiment spits out a random series of plus-ones and minus-ones. Determinism has broken down, but in a particular way. If we do many repetitions, we will find

that the numbers of $\sigma = +1$ events and $\sigma = -1$ events are statistically equal. In other words, the average value of σ is zero. Instead of the classical result—namely, that the component of σ along the x axis is zero—we find that the *average of these repeated measurements* is zero.

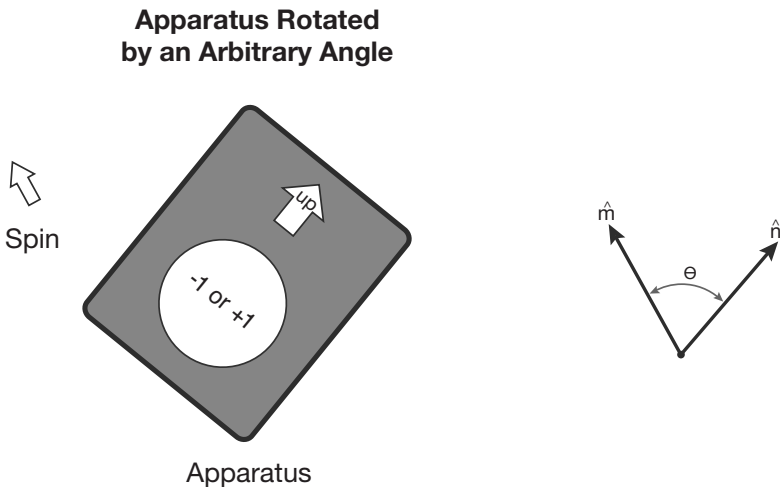


Figure 1.4: The apparatus rotated by an arbitrary angle within the x - z plane. Average measurement result is $\hat{n} \cdot \hat{m}$.

Now let's do the whole thing over again, but instead of rotating \mathcal{A} to lie on the x axis, rotate it to an arbitrary direction along the unit vector² \hat{n} . Classically, if σ were a vector, we would expect the result of the experiment to be the component of σ along the \hat{n} axis. If \hat{n} lies at an angle θ

²The standard notation for a unit vector (one of unit length) is to place a “hat” above the symbol representing the vector.

with respect to z , the classical answer would be $\sigma = \cos \theta$. But as you might guess, each time we do the experiment we get $\sigma = +1$ or $\sigma = -1$. However, the result is statistically biased so that the average value is $\cos \theta$.

The situation is of course more general. We did not have to start with \mathcal{A} oriented along z . Pick any direction \hat{m} and start with the up-arrow pointing along \hat{m} . Prepare a spin so that the apparatus reads $+1$. Then, without disturbing the spin, rotate the apparatus to the direction \hat{n} , as shown in Fig. 1.4. A new experiment on the same spin will give random results ± 1 , but with an average value equal to the cosine of the angle between \hat{n} and \hat{m} . In other words, the average will be $\hat{n} \cdot \hat{m}$.

The quantum mechanical notation for the statistical average of a quantity Q is Dirac's bracket notation $\langle Q \rangle$. We may summarize the results of our experimental investigation as follows: If we begin with \mathcal{A} oriented along \hat{m} and confirm that $\sigma = +1$, then subsequent measurement with \mathcal{A} oriented along \hat{n} gives the statistical result

$$\langle \sigma \rangle = \hat{n} \cdot \hat{m}.$$

What we are learning is that quantum mechanical systems are not deterministic—the results of experiments can be statistically random—but if we repeat an experiment many times, average quantities can follow the expectations of classical physics, at least up to a point.

1.4 Experiments Are Never Gentle

Every experiment involves an outside system—an apparatus—that must interact with the system in order to record a result. In that sense, every experiment is invasive. This is true in both classical and quantum physics, but only quantum physics makes a big deal out of it. Why is that so? Classically, an ideal measuring apparatus has a vanishingly small effect on the system it is measuring. Classical experiments can be arbitrarily gentle and still accurately and reproducibly record the results of the experiment. For example, the direction of an arrow can be determined by reflecting light off the arrow and focusing it to form an image. While it is true that the light must have a small enough wavelength to form an image, there is nothing in classical physics that prevents the image from being made with arbitrarily weak light. In other words, the light can have an arbitrarily small energy content.

In quantum mechanics, the situation is fundamentally different. *Any* interaction that is strong enough to measure some aspect of a system is necessarily strong enough to disrupt some other aspect of the same system. Thus, you can learn nothing about a quantum system without changing something else.

This should be evident in the examples involving \mathcal{A} and σ . Suppose we begin with $\sigma = +1$ along the z axis. If we measure σ again with \mathcal{A} oriented along z , we will confirm the previous value. We can do this over and over without changing the result. But consider this possibility: Between successive measurements along the z axis, we turn \mathcal{A} through

90 degrees, make an intermediate measurement, and turn it back to its original direction. Will a subsequent measurement along the z axis confirm the original measurement? The answer is no. The intermediate measurement along the x axis will leave the spin in a completely random configuration as far as the next measurement is concerned. There is no way to make the intermediate determination of the spin without completely disrupting the final measurement. One might say that measuring one component of the spin destroys the information about another component. In fact, one simply cannot simultaneously know the components of the spin along two different axes, not in a reproducible way in any case. There is something fundamentally different about the state of a quantum system and the state of a classical system.

1.5 Propositions

The space of states of a classical system is a mathematical set. If the system is a coin, the space of states is a set of two elements, H and T . Using set notation, we would write $\{H, T\}$. If the system is a six-sided die, the space of states has six elements labeled $\{1, 2, 3, 4, 5, 6\}$. The logic of set theory is called *Boolean* logic. Boolean logic is just a formalized version of the familiar classical logic of propositions.

A fundamental idea in Boolean logic is the notion of a truth-value. The truth-value of a proposition is either *true* or *false*. Nothing in between is allowed. The related set theory concept is a subset. Roughly speaking, a proposition is true for all the elements in its corresponding subset and false for all the elements not in this subset. For example,

if the set represents the possible states of a die, one can consider the proposition

A: The die shows an odd-numbered face.

The corresponding subset contains the three elements $\{1, 3, 5\}$.

Another proposition states

B: The die shows a number less than 4.

The corresponding subset contains the states $\{1, 2, 3\}$.

Every proposition has its opposite (also called its negation). For example,

not A: *The die does not show an odd-numbered face.*

The subset for this negated proposition is $\{2, 4, 6\}$.

There are rules for combining propositions into more complex propositions, the most important being **or**, **and**, and **not**. We just saw an example of **not**, which gets applied to a single subset or proposition. **And** is straightforward, and applies to a pair of propositions.³ It says they are both true. Applied to two subsets, **and** gives the elements common to both, that is, the *intersection* of the two subsets. In the die example, the intersection of subsets *A* and *B* is the subset of elements that are both odd and less than 4. Fig. 1.5 uses a Venn diagram to show how this works.

³**And** may be defined for multiple propositions, but we'll only consider two. The same goes for **or**.

The **or** rule is similar to **and**, but has one additional subtlety. In everyday speech, the word *or* is generally used in the exclusive sense—the exclusive version is true if one or the other of two propositions is true, but not both. However, Boolean logic uses the *inclusive* version of **or**, which is true if either or both of the propositions are true. Thus, according to the inclusive **or**, the proposition

Albert Einstein discovered relativity or Isaac Newton was English

is true. So is

Albert Einstein discovered relativity or Isaac Newton was Russian.

The inclusive **or** is only wrong if both propositions are false. For example,

Albert Einstein discovered America⁴ or Isaac Newton was Russian.

The inclusive **or** has a set theoretic interpretation as the union of two sets: it denotes the subset containing anything in either or both of the component subsets. In the die example, $(A \text{ or } B)$ denotes the subset $\{1, 2, 3, 5\}$.

⁴OK, perhaps Einstein did discover America. But he was not the first.

Space of States for a Single Die

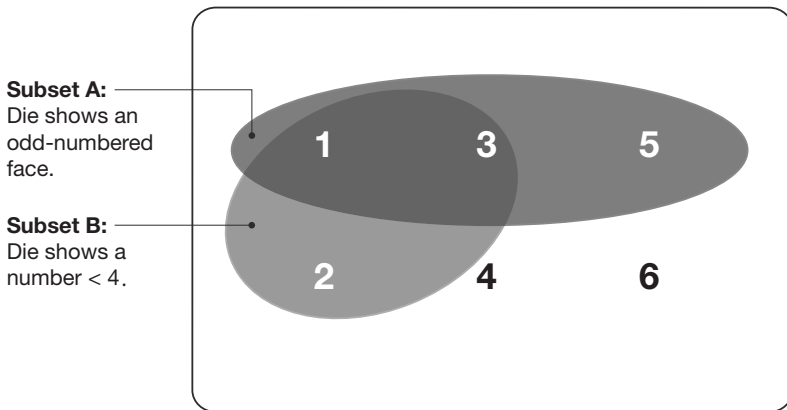


Figure 1.5: An Example of the Classical model of State Space. Subset A represents the proposition “the die shows an odd-numbered face.” Subset B : “The die shows a number < 4 .” Dark shading shows the intersection of A and B , which represents the proposition (A **and** B). White numbers are elements of the union of A with B , representing the proposition (A **or** B).

1.6 Testing Classical Propositions

Let’s return to the simple quantum system consisting of a single spin, and the various propositions whose truth we could test using the apparatus \mathcal{A} . Consider the following two propositions:

A : The z component of the spin is $+1$.

B : The x component of the spin is $+1$.

Each of these is meaningful and can be tested by orienting \mathcal{A} along the appropriate axis. The negation of each is also meaningful. For example, the negation of the first proposition is

not A: The z component of the spin is -1 .

But now consider the composite propositions

$(A \text{ or } B)$: The z component of the spin is $+1$ **or** the x component of the spin is $+1$.

$(A \text{ and } B)$: The z component of the spin is $+1$ **and** the x component of the spin is $+1$.

Consider how we would test the proposition $(A \text{ or } B)$. If spins behaved classically (and of course they don't), we would proceed as follows:⁵

- Gently measure σ_z and record the value. If it is $+1$, we are finished: the proposition $(A \text{ or } B)$ is true. If σ_z is -1 , continue to the next step.
- Gently measure σ_x . If it is $+1$, then the proposition $(A \text{ or } B)$ is true. If not, this means that neither σ_z nor σ_x was equal to $+1$, and $(A \text{ or } B)$ is false.

⁵Recall that the classical meaning of σ is different from the quantum mechanical meaning. Classically, σ is a straightforward 3-vector; σ_x and σ_z represent its spatial components.

There is an alternative procedure, which is to interchange the order of the two measurements. To emphasize this reversal of ordering, we'll call the new procedure (B **or** A):

- Gently measure σ_x and record the value. If it is $+1$ we are finished: The proposition (B **or** A) is true. If σ_x is -1 continue to the next step.
- Gently measure σ_z . If it is $+1$, then (B **or** A) is true. If not, it means that neither σ_x nor σ_z was equal to $+1$, and (B **or** A) is false.

In classical physics, the two orders of operation give the same answer. The reason for this is that measurements can be arbitrarily gentle—so gentle that they do not affect the results of subsequent measurements. Therefore, the proposition (A **or** B) has the same meaning as the proposition (B **or** A).

1.7 Testing Quantum Propositions

Now we come to the quantum world that I described earlier. Let us imagine a situation in which someone (or something) unknown to us has secretly prepared a spin in the $\sigma_z = +1$ state. Our job is to use the apparatus \mathcal{A} to determine whether the proposition (A **or** B) is true or false. We will try using the procedures outlined above.

We begin by measuring σ_z . Since the unknown agent has set things up, we will discover that $\sigma_z = +1$. It is unnecessary to go on: (A **or** B) is true. Nevertheless, we could test σ_x

just to see what happens. The answer is unpredictable. We randomly find that $\sigma_x = +1$ or $\sigma_x = -1$. But neither of these outcomes affects the truth of proposition (A **or** B).

But now let's reverse the order of measurement. As before, we'll call the reversed procedure (B **or** A), and this time we'll measure σ_x first. Because the unknown agent set the spin to $+1$ along the z axis, the measurement of σ_x is random. If it turns out that $\sigma_x = +1$, we are finished: (B **or** A) is true. But suppose we find the opposite result, $\sigma_x = -1$. The spin is oriented along the $-x$ direction. Let's pause here briefly, to make sure we understand what just happened. As a result of our first measurement, the spin is no longer in its original state $\sigma_z = +1$. It is in a new state, which is either $\sigma_x = +1$ or $\sigma_x = -1$. Please take a moment to let this idea sink in. We cannot overstate its importance.

Now we're ready to test the second half of proposition (B **or** A). Rotate the apparatus \mathcal{A} to the z axis and measure σ_z . According to quantum mechanics, the result will be randomly ± 1 . This means that there is a 25 percent probability that the experiment produces $\sigma_x = -1$ and $\sigma_z = -1$. In other words, with a probability of $\frac{1}{4}$, we find that (B **or** A) is false; this occurs despite the fact that the hidden agent had originally made sure that $\sigma_z = +1$.

Evidently, in this example, the inclusive **or** is not symmetric. The truth of (A **or** B) may depend on the order in which we confirm the two propositions. This is not a small thing; it means not only that the laws of quantum physics are different from their classical counterparts, but that the very foundations of logic are different in quantum physics as well.

What about (*A and B*)? Suppose our first measurement yields $\sigma_z = +1$ and the second, $\sigma_x = +1$. This is of course a possible outcome. We would be inclined to say that (*A and B*) is true. But in science, especially in physics, the truth of a proposition implies that the proposition can be verified by subsequent observation. In classical physics, the gentleness of observations implies that subsequent experiments are unaffected and will confirm an earlier experiment. A coin that turns up Heads will not be flipped to Tails by the act of observing it—at least not classically. Quantum mechanically, the second measurement ($\sigma_x = +1$) ruins the possibility of verifying the first. Once σ_x has been prepared along the x axis, another measurement of σ_z will give a random answer. Thus (*A and B*) is not confirmable: the second piece of the experiment interferes with the possibility of confirming the first piece.

If you know a bit about quantum mechanics, you probably recognize that we are talking about the uncertainty principle. The uncertainty principle doesn't apply only to position and momentum (or velocity); it applies to many pairs of measurable quantities. In the case of the spin, it applies to propositions involving two different components of σ . In the case of position and momentum, the two propositions we might consider are:

A certain particle has position x .

That same particle has momentum p .

From these, we can form the two composite propositions

The particle has position x **and** the particle has momentum p .

The particle has position x **or** the particle has momentum p .

Awkward as they are, both of these propositions have meaning in the English language, and in classical physics as well. However, in quantum physics, the first of these propositions is completely meaningless (not even wrong), and the second one means something quite different from what you might think. It all comes down to a deep logical difference between the classical and quantum concepts of the state of a system. Explaining the quantum concept of state will require some abstract mathematics, so let's pause for a brief interlude on complex numbers and vector spaces. The need for complex quantities will become clear later on, when we study the mathematical representation of spin states.

1.8 Mathematical Interlude: Complex Numbers

Everyone who has gotten this far in the Theoretical Minimum series knows about complex numbers. Nevertheless, I will spend a few lines reminding you of the essentials. Fig. 1.6 shows some of their basic elements.

A complex number z is the sum of a real number and an imaginary number. We can write it as

$$z = x + iy,$$

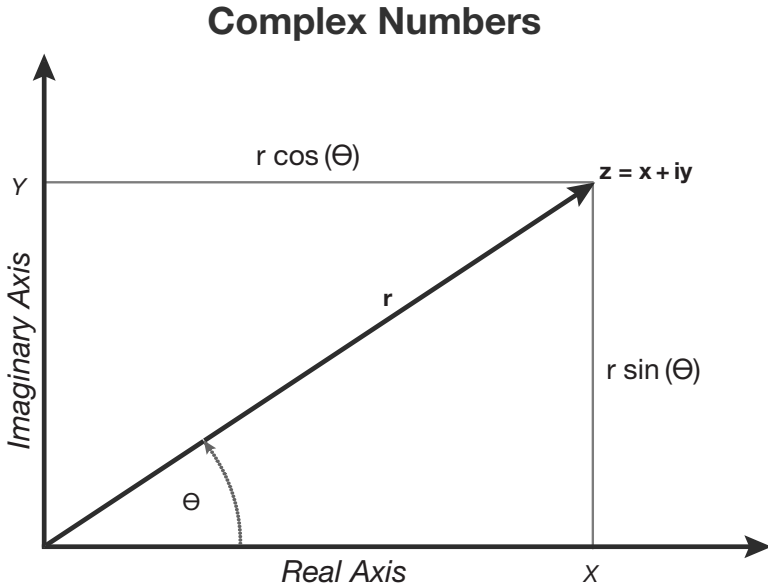


Figure 1.6: Two Common Ways to Represent Complex Numbers. In the Cartesian representation, x and y are the horizontal (real) and vertical (imaginary) components. In the polar representation, r is the radius, and θ is the angle made with the x axis. In each case, it takes two real numbers to represent a single complex number.

where x and y are real and $i^2 = -1$. Complex numbers can be added, multiplied, and divided by the standard rules of arithmetic. They can be visualized as points on the complex plane with coordinates x, y . They can also be represented in polar coordinates:

$$z = re^{i\theta} = r(\cos \theta + i \sin \theta).$$

Adding complex numbers is easy in component form: just add the components. Similarly, multiplying them is easy in their polar form: Simply multiply the radii and add the angles:

$$(r_1 e^{i\theta_1}) (r_2 e^{i\theta_2}) = (r_1 r_2) e^{i(\theta_1 + \theta_2)}.$$

Every complex number z has a complex conjugate z^* that is obtained by simply reversing the sign of the imaginary part. If

$$z = x + iy = re^{i\theta},$$

then

$$z^* = x - iy = re^{-i\theta}.$$

Multiplying a complex number and its conjugate always gives a positive real result:

$$z^* z = r^2.$$

It is of course true that every complex conjugate is itself a complex number, but it's often helpful to think of z and z^* as belonging to separate “dual” number systems. *Dual* here means that for every z there is a unique z^* and vice versa.

There is a special class of complex numbers that I'll call "phase-factors." A phase-factor is simply a complex number whose r -component is 1. If z is a phase-factor, then the following hold:

$$z^* z = 1$$

$$z = e^{i\theta}$$

$$z = \cos \theta + i \sin \theta.$$

1.9 Mathematical Interlude: Vector Spaces

1.9.1 Axioms

For a classical system, the space of states is a set (the set of possible states), and the logic of classical physics is Boolean. That seems obvious and it is difficult to imagine any other possibility. Nevertheless, the real world operates along entirely different lines, at least whenever quantum mechanics is important. The space of states of a quantum system is *not* a mathematical set;⁶ it is a *vector space*. Relations between the elements of a vector space are different from those between the elements of a set, and the logic of propositions is different as well.

Before I tell you about vector spaces, I need to clarify the term *vector*. As you know, we use this term to indicate an

⁶To be a little more precise, we will not focus on the set-theoretic properties of state spaces, even though they may of course be regarded as sets.

object in ordinary space that has a magnitude and a direction. Such vectors have three components, corresponding to the three dimensions of space. I want you to completely forget about that concept of a vector. From now on, whenever I want to talk about a thing with magnitude and direction in ordinary space, I will explicitly call it a *3-vector*. A mathematical vector space is an abstract construction that may or may not have anything to do with ordinary space. It may have any number of dimensions from 1 to ∞ and it may have components that are integers, real numbers, or even more general things.

The vector spaces we use to define quantum mechanical states are called *Hilbert spaces*. We won't give the mathematical definition here, but you may as well add this term to your vocabulary. When you come across the term *Hilbert space* in quantum mechanics, it refers to the space of states. A Hilbert space may have either a finite or an infinite number of dimensions.

In quantum mechanics, a vector space is composed of elements $|A\rangle$ called *ket-vectors* or just *kets*. Here are the axioms we will use to define the vector space of states of a quantum system (z and w are complex numbers):

1. The sum of any two ket-vectors is also a ket-vector:

$$|A\rangle + |B\rangle = |C\rangle.$$

2. Vector addition is commutative:

$$|A\rangle + |B\rangle = |B\rangle + |A\rangle.$$

3. Vector addition is associative:

$$\{|A\rangle + |B\rangle\} + |C\rangle = |A\rangle + \{|B\rangle + |C\rangle\}.$$

4. There is a unique vector 0 such that when you add it to any ket, it gives the same ket back:

$$|A\rangle + 0 = |A\rangle.$$

5. Given any ket $|A\rangle$, there is a unique ket $-|A\rangle$ such that

$$|A\rangle + (-|A\rangle) = 0.$$

6. Given any ket $|A\rangle$ and any complex number z , you can multiply them to get a new ket. Also, multiplication by a scalar is linear:

$$|zA\rangle = z|A\rangle = |B\rangle.$$

7. The distributive property holds:

$$\begin{aligned} z\{|A\rangle + |B\rangle\} &= z|A\rangle + z|B\rangle \\ \{z + w\}|A\rangle &= z|A\rangle + w|A\rangle. \end{aligned}$$

Axioms 6 and 7 taken together are often called *linearity*.

Ordinary 3-vectors would satisfy these axioms except for one thing: Axiom 6 allows a vector to be multiplied by any complex number. Ordinary 3-vectors can be multiplied by real numbers (positive, negative, or zero) but multiplication by complex numbers is not defined. One can think of 3-vectors as forming a real vector space, and kets as forming a complex vector space. Our definition of ket-vectors is fairly abstract. As we will see, there are various concrete ways to represent ket-vectors as well.

1.9.2 Functions and Column Vectors

Let's look at some concrete examples of complex vector spaces. First of all, consider the set of continuous complex-valued functions of a variable x . Call the functions $A(x)$. You can add any two such functions and multiply them by complex numbers. You can check that they satisfy all seven axioms. This example should make it obvious that we are talking about something much more general than three-dimensional arrows.

Two-dimensional column vectors provide another concrete example. We construct them by stacking up a pair of complex numbers, α_1 and α_2 , in the form

$$\begin{pmatrix} \alpha_1 \\ \alpha_2 \end{pmatrix}$$

and identifying this “stack” with the ket-vector $|A\rangle$. The complex numbers α are the components of $|A\rangle$. You can add two column vectors by adding their components:

$$\begin{pmatrix} \alpha_1 \\ \alpha_2 \end{pmatrix} + \begin{pmatrix} \beta_1 \\ \beta_2 \end{pmatrix} = \begin{pmatrix} \alpha_1 + \beta_1 \\ \alpha_2 + \beta_2 \end{pmatrix}.$$

Moreover, you can multiply the column vector by a complex number z just by multiplying the components,

$$z \begin{pmatrix} \alpha_1 \\ \alpha_2 \end{pmatrix} = \begin{pmatrix} z\alpha_1 \\ z\alpha_2 \end{pmatrix}.$$

Column vector spaces of any number of dimensions can be constructed. For example, here is a five-dimensional column vector:

$$\begin{pmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_3 \\ \alpha_4 \\ \alpha_5 \end{pmatrix}.$$

Normally, we do not mix vectors of different dimensionality.

1.9.3 Bras and Kets

As we have seen, the complex numbers have a dual version: in the form of complex conjugate numbers. In the same way, a complex vector space has a dual version that is essentially the complex conjugate vector space. For every ket-vector $|A\rangle$, there is a “bra” vector in the dual space, denoted by $\langle A|$. Why the strange terms *bra* and *ket*? Shortly, we will define inner products of bras and kets, using expressions like $\langle B|A\rangle$ to form *bra-kets* or brackets. Inner products are extremely

important in the mathematical machinery of quantum mechanics, and for characterizing vector spaces in general.

Bra vectors satisfy the same axioms as the ket-vectors, but there are two things to keep in mind about the correspondence between kets and bras:

1. Suppose $\langle A|$ is the bra corresponding to the ket $|A\rangle$, and $\langle B|$ is the bra corresponding to the ket $|B\rangle$. Then the bra corresponding to

$$|A\rangle + |B\rangle$$

is

$$\langle A| + \langle B|.$$

2. If z is a complex number, then it is *not* true that the bra corresponding to $z|A\rangle$ is $\langle A|z$. You have to remember to complex-conjugate. Thus, the bra corresponding to

$$z|A\rangle$$

is

$$\langle A|z^*.$$

In the concrete example where kets are represented by column vectors, the dual bras are represented by row vectors,

with the entries being drawn from the complex conjugate numbers. Thus, if the ket $|A\rangle$ is represented by the column

$$\begin{pmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_3 \\ \alpha_4 \\ \alpha_5 \end{pmatrix},$$

then the corresponding bra $\langle A|$ is represented by the row

$$\left(\alpha_1^* \quad \alpha_2^* \quad \alpha_3^* \quad \alpha_4^* \quad \alpha_5^* \right).$$

1.9.4 Inner Products

You are no doubt familiar with the dot product defined for ordinary 3-vectors. The analogous operation for bras and kets is the *inner product*. The inner product is always the product of a bra and a ket and it is written this way:

$$\langle B|A\rangle.$$

The result of this operation is a complex number. The axioms for inner products are not too hard to guess:

1. They are linear:

$$\langle C| \{ |A\rangle + |B\rangle \} = \langle C|A\rangle + \langle C|B\rangle.$$

2. Interchanging bras and kets corresponds to complex conjugation:

$$\langle B|A \rangle = \langle A|B \rangle^*.$$

Exercise 1.1:

- a) Using the axioms for inner products, prove

$$\{ \langle \mathbf{A} | + \langle \mathbf{B} | \} | \mathbf{C} \rangle = \langle \mathbf{A} | \mathbf{C} \rangle + \langle \mathbf{B} | \mathbf{C} \rangle.$$

- b) Prove $\langle \mathbf{A} | \mathbf{A} \rangle$ is a real number.

In the concrete representation of bras and kets by row and column vectors, the inner product is defined in terms of components:

$$\begin{aligned} \langle B|A \rangle &= \left(\beta_1^* \quad \beta_2^* \quad \beta_3^* \quad \beta_4^* \quad \beta_5^* \right) \begin{pmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_3 \\ \alpha_4 \\ \alpha_5 \end{pmatrix} \\ &= \beta_1^* \alpha_1 + \beta_2^* \alpha_2 + \beta_3^* \alpha_3 + \beta_4^* \alpha_4 + \beta_5^* \alpha_5. \quad (1.2) \end{aligned}$$

The rule for inner products is essentially the same as for dot products: add the products of corresponding components of the vectors whose inner product is being calculated.

Exercise 1.2: Show that the inner product defined by Eq. 1.2 satisfies all the axioms of inner products.

Using the inner product, we can define some concepts that are familiar from ordinary 3-vectors:

- *Normalized Vector:* A vector is said to be normalized if its inner product with itself is 1. Normalized vectors satisfy,

$$\langle A|A \rangle = 1.$$

For ordinary 3-vectors, the term *normalized vector* is usually replaced by *unit vector*, that is, a vector of unit length.

- *Orthogonal Vectors:* Two vectors are said to be orthogonal if their inner product is zero. $|A\rangle$ and $|B\rangle$ are orthogonal if

$$\langle B|A \rangle = 0.$$

This is the analog of saying that two 3-vectors are orthogonal if their dot product is zero.

1.9.5 Orthonormal Bases

When working with ordinary 3-vectors, it is extremely useful to introduce a set of three mutually orthogonal unit vectors and use them as a basis to construct any vector. A simple

example would be the unit 3-vectors that point along the x , y , and z axes. They are usually called \hat{i} , \hat{j} , and \hat{k} . Each is of unit length and orthogonal to the others. If you tried to find a fourth vector orthogonal to these three, there wouldn't be any—not in three dimensions anyway. However, if there were more dimensions of space, there would be more basis vectors. The dimension of a space can be defined as the maximum number of mutually orthogonal vectors in that space.

Obviously, there is nothing special about the particular axes x , y , and z . As long as the basis vectors are of unit length and are mutually orthogonal, they comprise an *orthonormal basis*.

The same principle is true for complex vector spaces. One can begin with any normalized vector and then look for a second one, orthogonal to the first. If you find one, then the space is at least two-dimensional. Then look for a third, fourth, and so on. Eventually, you may run out of new directions and there will not be any more orthogonal candidates. The maximum number of mutually orthogonal vectors is the dimension of the space. For column vectors, the dimension is simply the number of entries in the column.

Let's consider an N -dimensional space and a particular orthonormal basis of ket-vectors labeled $|i\rangle$.⁷ The label i runs from 1 to N . Consider a vector $|A\rangle$, written as a sum of basis

⁷Mathematically, basis vectors are not required to be orthonormal. However, in quantum mechanics they generally are. In this book, whenever we say *basis*, we mean an orthonormal basis.

vectors:

$$|A\rangle = \sum_i \alpha_i |i\rangle. \quad (1.3)$$

The α_i are complex numbers called the *components* of the vector, and to calculate them we take the inner product of both sides with a basis bra $\langle j|$:

$$\langle j|A\rangle = \sum_i \alpha_i \langle j|i\rangle. \quad (1.4)$$

Next, we use the fact that the basis vectors are orthonormal. This implies that $\langle j|i\rangle = 0$ if i is not equal to j , and $\langle j|i\rangle = 1$ if $i = j$. In other words, $\langle j|i\rangle = \delta_{ij}$. This makes the sum in Eq. 1.4 collapse to one term:

$$\langle j|A\rangle = \alpha_j. \quad (1.5)$$

Thus, we see that the components of a vector are just its inner products with the basis vectors. We can rewrite Eq. 1.3 in the elegant form

$$|A\rangle = \sum_i |i\rangle \langle i|A\rangle.$$

Lecture 2

Quantum States

Art: *Oddly enough, that beer made my head **stop** spinning. What state are we in?*

Lenny: *I wish I knew. Does it matter?*

Art: *It might. I don't think we're in California anymore.*

2.1 States and Vectors

In classical physics, knowing *the state of a system* implies knowing everything that is necessary to predict the future of that system. As we've seen in the last lecture, quantum systems are not completely predictable. Evidently, quantum states have a different meaning than classical states. Very roughly, knowing a quantum state means knowing *as much as can be known* about how the system was prepared. In the last chapter, we talked about using an apparatus to prepare the state of a spin. In fact, we implicitly assumed that there

was no more fine detail to specify or that could be specified about the state of the spin.

The obvious question to ask is whether the unpredictability is due to an incompleteness in what we call a quantum state. There are various opinions about this matter. Here is a sampling:

- Yes, the usual notion of quantum state is incomplete. There are “hidden variables” that, if only we could access them, would allow complete predictability. There are two versions of this view. In version *A*, the hidden variables are hard to measure but in principle they are experimentally available to us. In version *B*, because we are made of quantum mechanical matter and therefore subject to the restrictions of quantum mechanics, the hidden variables are, in principle, not detectable.
- No, the hidden variables concept does not lead us in a profitable direction. Quantum mechanics is unavoidably unpredictable. Quantum mechanics is as complete a calculus of probabilities as is possible. The job of a physicist is to learn and use this calculus.

I don't know what the ultimate answer to this question will be, or even if it will prove to be a useful question. But for our purposes, it's not important what any particular physicist believes about the ultimate meaning of the quantum state. For practical reasons, we will adopt the second view.

In practice, what this means for the quantum spin of Lecture 1 is that, when the apparatus \mathcal{A} acts and tells us that $\sigma_z = +1$ or $\sigma_z = -1$, there is no more to know, or

that can be known. Likewise, if we rotate \mathcal{A} and measure $\sigma_x = +1$ or $\sigma_x = -1$, there is no more to know. Likewise for σ_y or any other component of the spin.

2.2 Representing Spin States

Now it's time to try our hand at representing spin states using state-vectors. Our goal is to build a representation that captures everything we know about the behavior of spins. At this point, the process will be more intuitive than formal. We will try to fit things together the best we can, based on what we've already learned. Please read this section carefully. Believe me, it will pay off.

Let's begin by labeling the possible spin states along the three coordinate axes. If \mathcal{A} is oriented along the z axis, the two possible states that can be prepared correspond to $\sigma_z = \pm 1$. Let's call them *up* and *down* and denote them by ket-vectors $|u\rangle$ and $|d\rangle$. Thus, when the apparatus is oriented along the z axis and registers $+1$, the state $|u\rangle$ has been prepared.

On the other hand, if the apparatus is oriented along the x axis and registers -1 , the state $|l\rangle$ has been prepared. We'll call it *left*. If \mathcal{A} is along the y axis, it can prepare the states $|i\rangle$ and $|o\rangle$ (*in* and *out*). You get the idea.

The idea that there are no hidden variables has a very simple mathematical representation: the space of states for a single spin has only two dimensions. This point deserves emphasis:

All possible spin states can be represented in a two-dimensional vector space.

We could, somewhat arbitrarily,¹ choose $|u\rangle$ and $|d\rangle$ as the two basis vectors and write *any* state as a linear superposition of these two. We'll adopt that choice for now. Let's use the symbol $|A\rangle$ for a generic state. We can write this as an equation,

$$|A\rangle = \alpha_u|u\rangle + \alpha_d|d\rangle,$$

where α_u and α_d are the components of $|A\rangle$ along the basis directions $|u\rangle$ and $|d\rangle$. Mathematically, we can identify the components of $|A\rangle$ as

$$\begin{aligned}\alpha_u &= \langle u|A\rangle \\ \alpha_d &= \langle d|A\rangle.\end{aligned}\tag{2.1}$$

These equations are extremely abstract, and it is not at all obvious what their physical significance is. I am going to tell you right now what they mean: First of all, $|A\rangle$ can represent any state of the spin, prepared in any manner. The components α_u and α_d are complex numbers; by themselves, they have no experimental meaning, but their magnitudes do. In particular, $\alpha_u^*\alpha_u$ and $\alpha_d^*\alpha_d$ have the following meaning:

- Given that the spin has been prepared in the state $|A\rangle$, and that the apparatus is oriented along z , the

¹The choice is not totally arbitrary. The basis vectors must be orthogonal to each other.

quantity $\alpha_u^* \alpha_u$ is the probability that the spin would be measured as $\sigma_z = +1$. In other words, it is the probability of the spin being *up* if measured along the z axis.

- Likewise, $\alpha_d^* \alpha_d$ is the probability that σ_z would be *down* if measured.

The α values, or equivalently $\langle u|A\rangle$ and $\langle d|A\rangle$, are called probability amplitudes. They are themselves not probabilities. To compute a probability, their magnitudes must be squared. In other words, the probabilities for measurements of *up* and *down* are given by

$$\begin{aligned} P_u &= \langle A|u\rangle\langle u|A\rangle \\ P_d &= \langle A|d\rangle\langle d|A\rangle. \end{aligned} \tag{2.2}$$

Notice that I have said nothing about what σ_z is before it is measured. Before the measurement, all we have is the vector $|A\rangle$, which represents the potential possibilities but not the actual values of our measurements.

Two other points are important: First, note that $|u\rangle$ and $|d\rangle$ are mutually orthogonal. In other words,

$$\begin{aligned} \langle u|d\rangle &= 0 \\ \langle d|u\rangle &= 0. \end{aligned} \tag{2.3}$$

The physical meaning of this is that, if the spin is prepared *up*, then the probability to detect it *down* is zero, and vice versa. This point is so important, I'll say it again: Two orthogonal states are physically distinct and mutually exclusive. If the spin is in one of these states, it *cannot* be (has zero probability to be) in the other one. This idea applies to all quantum systems, not just spin.

But don't mistake the orthogonality of state-vectors for orthogonal directions in space. In fact, the directions *up* and *down* are not orthogonal directions in space, even though their associated state-vectors are orthogonal in state space.

The second important point is that for the total probability to come out equal to unity, we must have

$$\alpha_u^* \alpha_u + \alpha_d^* \alpha_d = 1. \quad (2.4)$$

This is equivalent to saying that the vector $|A\rangle$ is normalized to a unit vector:

$$\langle A|A\rangle = 1.$$

This is a very general principle of quantum mechanics that extends to all quantum systems: the state of a system is represented by a unit (normalized) vector in a vector space of states. Moreover, the squared magnitudes of the components of the state-vector, along particular basis vectors, represent probabilities for various experimental outcomes.

2.3 Along the x Axis

We said before that we can represent any spin state as a linear combination of the basis vectors $|u\rangle$ and $|d\rangle$. Let's try doing this now for the vectors $|r\rangle$ and $|l\rangle$, which represent spins prepared along the x axis. We'll start with $|r\rangle$. As you recall from Lecture 1, if \mathcal{A} initially prepares $|r\rangle$, and is then rotated to measure σ_z , there will be equal probabilities for *up* and *down*. Thus, $\alpha_u^* \alpha_u$ and $\alpha_d^* \alpha_d$ must both be equal to $\frac{1}{2}$. A simple vector that satisfies this rule is

$$|r\rangle = \frac{1}{\sqrt{2}}|u\rangle + \frac{1}{\sqrt{2}}|d\rangle. \quad (2.5)$$

There is some ambiguity in this choice, but as we will see later, it is nothing more than the ambiguity in our choice of exact directions for the x and y axes.

Next, let's look at the vector $|l\rangle$. Here is what we know: when the spin has been prepared in the *left* configuration, the probabilities for σ_z are again equal to $\frac{1}{2}$. That is not enough to determine the values $\alpha_u^* \alpha_u$ and $\alpha_d^* \alpha_d$, but there is another condition that we can infer. Earlier, I told you that $|u\rangle$ and $|d\rangle$ are orthogonal for the simple reason that, if the spin is *up*, it's definitely not *down*. But there is nothing special about *up* and *down* that is not also true of *right* and *left*. In particular, if the spin is *right*, it has zero probability of being *left*. Thus, by analogy with Eq. 2.3,

$$\langle r|l\rangle = 0$$

$$\langle l|r\rangle = 0.$$

This pretty much fixes $|l\rangle$ in the form

$$|l\rangle = \frac{1}{\sqrt{2}}|u\rangle - \frac{1}{\sqrt{2}}|d\rangle. \quad (2.6)$$

Exercise 2.1: Prove that the vector $|r\rangle$ in Eq. 2.5 is orthogonal to vector $|l\rangle$ in Eq. 2.6.

Again, there is some ambiguity in the choice of $|l\rangle$. This is called the *phase ambiguity*. Suppose we multiply $|l\rangle$ by any complex number z . That will have no effect on whether it is orthogonal to $|r\rangle$, though in general the result will no longer be normalized (have unit length). But if we choose $z = e^{i\theta}$ (where θ can be any real number), then there will be no effect on the normalization because $e^{i\theta}$ has unit magnitude. In other words, $\alpha_u^* \alpha_u + \alpha_d^* \alpha_d$ will remain equal to 1. Since a number of the form $z = e^{i\theta}$ is called a phase-factor, the ambiguity is called the phase ambiguity. Later, we will find out that no measurable quantity is sensitive to the overall phase-factor, and therefore we can ignore it when specifying states.

2.4 Along the y Axis

Finally, this brings us to $|i\rangle$ and $|o\rangle$, the vectors representing spins oriented along the y axis. Let's look at the conditions they need to satisfy. First,

$$\langle i|o\rangle = 0. \quad (2.7)$$

This condition states that *in* and *out* are represented by orthogonal vectors in the same way that *up* and *down* are. Physically, this means that if the spin is *in*, it is definitely not *out*.

There are additional restrictions on the vectors $|i\rangle$ and $|o\rangle$. Using the relationships expressed in Eqs. 2.1 and 2.2, and the statistical results of our experiments, we can write the following:

$$\begin{aligned}
 \langle o|u\rangle\langle u|o\rangle &= \frac{1}{2} \\
 \langle o|d\rangle\langle d|o\rangle &= \frac{1}{2} \\
 \langle i|u\rangle\langle u|i\rangle &= \frac{1}{2} \\
 \langle i|d\rangle\langle d|i\rangle &= \frac{1}{2}.
 \end{aligned} \tag{2.8}$$

In the first two equations, $|o\rangle$ takes the role of $|A\rangle$ from Eqs. 2.1 and 2.2. In the second two, $|i\rangle$ takes that role. These conditions state that if the spin is oriented along y , and is then measured along z , it is equally likely to be *up* or *down*. We should also expect that if the spin were measured along the x axis, it would be equally likely to be *right* or *left*. This leads to additional conditions:

$$\begin{aligned}
 \langle o|r\rangle\langle r|o\rangle &= \frac{1}{2} \\
 \langle o|l\rangle\langle l|o\rangle &= \frac{1}{2}
 \end{aligned}$$

$$\begin{aligned}\langle i|r\rangle\langle r|i\rangle &= \frac{1}{2} \\ \langle i|l\rangle\langle l|i\rangle &= \frac{1}{2}.\end{aligned}\tag{2.9}$$

These conditions are sufficient to determine the form of the vectors $|i\rangle$ and $|o\rangle$, apart from the phase ambiguity. Here is the result:

$$\begin{aligned}|i\rangle &= \frac{1}{\sqrt{2}}|u\rangle + \frac{i}{\sqrt{2}}|d\rangle \\ |o\rangle &= \frac{1}{\sqrt{2}}|u\rangle - \frac{i}{\sqrt{2}}|d\rangle.\end{aligned}\tag{2.10}$$

Exercise 2.2: Prove that $|i\rangle$ and $|o\rangle$ satisfy all of the conditions in Eqs. 2.7, 2.8, and 2.9. Are they unique in that respect?

It's interesting that two of the components in Eqs. 2.10 are imaginary. Of course, we've said all along that the space of states is a complex vector space, but until now we have not had to use complex numbers in our calculations. Are the complex numbers in Eqs. 2.10 a convenience or a necessity? Given our framework for spin states, there is no way around them. It's somewhat tedious to demonstrate this, but the steps are straightforward. The following exercise gives you a road map. The need for complex numbers is a general feature of quantum mechanics, and we'll see more examples as we go.

Exercise 2.3: For the moment, forget that Eqs. 2.10 give us working definitions for $|i\rangle$ and $|o\rangle$ in terms of $|u\rangle$ and $|d\rangle$, and assume that the components α , β , γ , and δ are unknown:

$$|i\rangle = \alpha|u\rangle + \beta|d\rangle$$

$$|o\rangle = \gamma|u\rangle + \delta|d\rangle.$$

a) Use Eqs. 2.8 to show that

$$\alpha^*\alpha = \beta^*\beta = \gamma^*\gamma = \delta^*\delta = \frac{1}{2}.$$

b) Use the above result and Eqs. 2.9 to show that

$$\alpha^*\beta + \alpha\beta^* = \gamma^*\delta + \gamma\delta^* = 0.$$

c) Show that $\alpha^*\beta$ and $\gamma^*\delta$ must each be pure imaginary.

If $\alpha^*\beta$ is pure imaginary, then α and β cannot both be real. The same reasoning applies to $\gamma^*\delta$.

2.5 Counting Parameters

It's always important to know how many independent parameters it takes to characterize a system. For example, the generalized coordinates we used in *Volume I* (referred to as q_i) each represented an independent degree of freedom. That approach freed us from the difficult job of writing explicit equations to describe physical constraints. Along similar lines, our next task is to count the number of physically distinct states there are for a spin. I will do it in two ways, to show that you get the same answer either way.

The first way is simple. Point the apparatus along any unit 3-vector² \hat{n} and prepare a spin with $\sigma = +1$ along that axis. If $\sigma = -1$, you can think of the spin as being oriented along the $-\hat{n}$ axis. Thus, there must be a state for every orientation of the unit 3-vector \hat{n} . How many parameters does it take to specify such an orientation? The answer is of course two. It takes two angles to define a direction in three-dimensional space.³

Now, let's consider the same question from another perspective. The general spin state is defined by two complex numbers, α_u and α_d . That seems to add up to four real parameters, with each complex parameter counting as two real ones. But recall that the vector has to be normalized as in Eq. 2.4. The normalization condition gives us one equation involving real variables, and cuts the number of parameters down to three.

As I said earlier, we will eventually see that the physical properties of a state-vector do not depend on the overall phase-factor. This means that one of the three remaining parameters is redundant, leaving only two—the same as the number of parameters we need to specify a direction in three-dimensional space. Thus, there is enough freedom in the expression

$$\alpha_u|u\rangle + \alpha_d|d\rangle$$

²Keep in mind that 3-vectors are not bras or kets.

³Recall that spherical coordinates use two angles to represent the orientation of a point in relation to the origin. Latitude and longitude provide another example.

to describe all the possible orientations of a spin, even though there are only two possible outcomes of an experiment along any axis.

2.6 Representing Spin States as Column Vectors

So far, we have been able to learn a lot by using the abstract forms of our state-vectors, that is, $|u\rangle$ and $|d\rangle$ and so forth. These abstractions help us focus on mathematical relationships without worrying about unnecessary details. However, soon we will need to perform detailed calculations on spin states, and for that we'll need to write our state-vectors in column form. Because of "phase indifference," the column representations are not unique, and we'll try to choose the simplest and most convenient ones we can find.

As usual, we'll start with $|u\rangle$ and $|d\rangle$. We need them to have unit length, and to be mutually orthogonal. A pair of columns that satisfies these requirements is

$$|u\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \quad (2.11)$$

$$|d\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}. \quad (2.12)$$

With these column vectors in hand, it will be easy to create column vectors for $|r\rangle$ and $|l\rangle$ using Eqs. 2.5 and 2.6, and for $|i\rangle$ and $|o\rangle$ using Eqs. 2.10. We'll do that in the next lecture, where these results are needed.

2.7 Putting It All Together

We have covered a lot of ground in this lecture. Before moving on, let's take stock of what we've done. Our goal was to synthesize what we know about spins and vector spaces. We figured out how to use vectors to represent spin states, and in the process we got a glimpse of the kind of information a state-vector contains (and does *not* contain!). Here is a brief outline of what we did:

- Based on our knowledge of spin measurements, we chose three pairs of mutually orthogonal basis vectors. Pair-wise, we named them $|u\rangle$ and $|d\rangle$, $|r\rangle$ and $|l\rangle$, and $|i\rangle$ and $|o\rangle$. Because the basis vectors $|u\rangle$ and $|d\rangle$ represent physically distinct states, we were able to assert that they are mutually orthogonal. In other words, $\langle u|d\rangle = 0$. The same holds for $|r\rangle$ and $|l\rangle$, and also for $|i\rangle$ and $|o\rangle$.
- We found that it takes two independent parameters to specify a spin state, and then we arbitrarily chose one of the orthogonal pairs, $|u\rangle$ and $|d\rangle$, as our basis vectors for representing *all* spin states—even though the two complex numbers in a state-vector require four real numbers to specify them. How did we get away with this? We were clever enough to notice that these four numbers are not all independent.⁴ The normalization constraint (total probability must equal 1) eliminates one independent parameter, and “phase indifference”

⁴Please indulge in a self-satisfied grin.

(the physics of a state-vector is unaffected by its overall phase-factor) eliminates a second.

- Having chosen $|u\rangle$ and $|d\rangle$ as our main basis vectors, we figured out how to represent the other two pairs of basis vectors as linear combinations of $|u\rangle$ and $|d\rangle$, using additional orthogonality and probability-based constraints.
- Finally, we established a way to represent our main basis vectors as columns. This representation is not unique. In the next lecture, we'll use our $|u\rangle$ and $|d\rangle$ column vectors to derive column vectors for the two other bases.

While achieving these concrete results, we got a chance to see some state-vector mathematics in action and learn something about how these mathematical objects correspond to physical spins. Although we will focus on spin, the same concepts and techniques apply to other quantum systems as well. Please take a little time to assimilate the material we've covered so far before moving on to the next lecture. As I said at the beginning, it will really pay off.

Lecture 3

Principles of Quantum Mechanics

Art: *I'm not like you, Lenny. My brain just wasn't built for quantum mechanics.*

Lenny: *Nah, mine wasn't either. Just can't really visualize the stuff. But I'll tell you, I once knew a guy who thought just like an electron.*

Art: *What happened to him?*

Lenny: *Art, all I'm gonna tell you is that it sure wasn't pretty.*

Art: *Hmm, I guess that gene didn't fly.*

No, we were not built to sense quantum phenomena; not the same way we were built to sense classical things like force and temperature. But we are very adaptable creatures and we've been able to substitute abstract mathematics for

the missing senses that might have allowed us to directly visualize quantum mechanics. And eventually we do develop new kinds of intuition.

This lecture introduces the principles of quantum mechanics. In order to describe those principles, we'll need some new mathematical tools. Let's get started.

3.1 Mathematical Interlude: Linear Operators

3.1.1 Machines and Matrices

States in quantum mechanics are mathematically described as vectors in a vector space. Physical observables—the things that you can measure—are described by linear operators. We'll take that as an axiom, and we'll find out later (in Section 3.1.5) that operators corresponding to physical observables must be Hermitian as well as linear. The correspondence between operators and observables is subtle, and understanding it will take some effort.

Observables are the things you measure. For example, we can make direct measurements of the coordinates of a particle; the energy, momentum, or angular momentum of a system; or the electric field at a point in space. Observables are also associated with a vector space, but they are not state-vectors. They are the things you measure— σ_x would be an example—and they are represented by *linear operators*. John Wheeler liked to call such mathematical objects *machines*. He imagined a machine with two ports: an input

port and an output port. In the input port you insert a vector, such as $|A\rangle$. The gears turn and the machine delivers a result in the output port. This result is another vector, say $|B\rangle$.

Let's denote the operator by the boldface letter \mathbf{M} (for "machine"). Here is the equation to express the fact that \mathbf{M} acts on the vector $|A\rangle$ to give $|B\rangle$:

$$\mathbf{M}|A\rangle = |B\rangle.$$

Not every machine is a *linear* operator. Linearity implies a few simple properties. To begin with, a linear operator must give a unique output for every vector in the space. We can imagine a machine that gives an output for some vectors, but just grinds up others and gives nothing. This machine would not be a linear operator. Something must come out for anything you put in.

The next property states that when a linear operator \mathbf{M} acts on a multiple of an input vector, it gives the same multiple of the output vector. Thus, if $\mathbf{M}|A\rangle = |B\rangle$, and z is any complex number, then

$$\mathbf{M}z|A\rangle = z|B\rangle.$$

The only other rule is that, when \mathbf{M} acts on a sum of vectors, the results are simply added together:

$$\mathbf{M}\{|A\rangle + |B\rangle\} = \mathbf{M}|A\rangle + \mathbf{M}|B\rangle.$$

To give a concrete representation of linear operators, we return to the row and column vector representation of bra-

and ket-vectors that we used in Lecture 1. The row-column notation depends on our choice of basis vectors. If the vector space is N -dimensional, we choose a set of N orthonormal (orthogonal and normalized) ket-vectors. Let's label them $|j\rangle$, and their dual bra-vectors $\langle j|$.

We are now going to take the equation

$$\mathbf{M}|A\rangle = |B\rangle$$

and write it in component form. As we did in Eq. 1.3, we'll represent an arbitrary ket $|A\rangle$ as a sum over basis vectors:

$$|A\rangle = \sum_j \alpha_j |j\rangle.$$

Here, we're using j as an index rather than i so you won't be tempted to think that we're talking about the *in* spin state. Now, we'll represent $|B\rangle$ in the same way and plug both of these substitutions into $\mathbf{M}|A\rangle = |B\rangle$. That gives

$$\sum_j \mathbf{M}|j\rangle \alpha_j = \sum_j \beta_j |j\rangle.$$

The last step is to take the inner product of both sides with a particular basis vector $\langle k|$, resulting in

$$\sum_j \langle k|\mathbf{M}|j\rangle \alpha_j = \sum_j \beta_j \langle k|j\rangle. \quad (3.1)$$

To make sense of this result, remember that $\langle k|j\rangle$ is zero if j and k are not equal, and 1 if they are equal. That means that the sum on the right side collapses to a single term, β_k .

On the left side, we see a set of quantities $\langle k|\mathbf{M}|j\rangle\alpha_j$. We can abbreviate $\langle k|\mathbf{M}|j\rangle$ with the symbol m_{kj} . Notice that each m_{kj} is just a complex number. To see why, think of \mathbf{M} operating on $|j\rangle$ to give some new ket-vector. The inner product of $\langle k|$ with this new ket-vector must be a complex number. The quantities m_{kj} are called the *matrix elements* of \mathbf{M} and are often arranged into a square $N \times N$ matrix. For example, if $N = 3$, we can write the symbolic equation

$$\mathbf{M} = \begin{pmatrix} m_{11} & m_{12} & m_{13} \\ m_{21} & m_{22} & m_{23} \\ m_{31} & m_{32} & m_{33} \end{pmatrix}. \quad (3.2)$$

This equation involves a slight abuse of notation that would give a purist indigestion. The left side is an abstract linear operator and the right side is a concrete representation of it in a particular basis. Equating them is sloppy but it should not cause confusion.

Now let's revisit Eq. 3.1 and replace $\langle k|\mathbf{M}|j\rangle$ with m_{kj} . We get

$$\sum_j m_{kj}\alpha_j = \beta_k. \quad (3.3)$$

We can write this in matrix form as well. Eq. 3.3 becomes

$$\begin{pmatrix} m_{11} & m_{12} & m_{13} \\ m_{21} & m_{22} & m_{23} \\ m_{31} & m_{32} & m_{33} \end{pmatrix} \begin{pmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_3 \end{pmatrix} = \begin{pmatrix} \beta_1 \\ \beta_2 \\ \beta_3 \end{pmatrix}. \quad (3.4)$$

You're probably familiar with the rule for matrix multiplication, but I will remind you just in case. To compute the first entry on the right, β_1 , take the first row of the matrix and "dot" it into the α column:

$$\beta_1 = m_{11}\alpha_1 + m_{12}\alpha_2 + m_{13}\alpha_3.$$

For the second entry, dot the second row of the matrix with the α column:

$$\beta_2 = m_{21}\alpha_1 + m_{22}\alpha_2 + m_{23}\alpha_3.$$

And so on. If you are not familiar with matrix multiplication, run to your computer and look it up right away. It's a crucial part of our tool kit, and I will assume you know it from now on.

There are both advantages and disadvantages to representing vectors and linear operators concretely with columns, rows, and matrices (known collectively as *components*). The advantages are obvious. Components provide a completely explicit set of arithmetic rules for working the machine. The disadvantage is that they depend on a specific choice of basis vectors. The underlying relationships between vectors and operators is independent of the particular basis we choose, and the concrete representation obscures that fact.

3.1.2 Eigenvalues and Eigenvectors

In general, when a linear operator acts on a vector, it will change the direction of the vector. This means that what

comes out of the machine will not just be the input vector multiplied by a number. But for a particular linear operator, there will be certain vectors whose directions are the same when they come out as they were when they went in. These special vectors are called *eigenvectors*. The definition of an eigenvector of \mathbf{M} is a vector $|\lambda\rangle$ such that

$$\mathbf{M}|\lambda\rangle = \lambda|\lambda\rangle. \quad (3.5)$$

The double use of λ is admittedly a little confusing. First of all, λ (as opposed to $|\lambda\rangle$) is a number—generally a complex one, but still a number. On the other hand, $|\lambda\rangle$ is a ket-vector. Furthermore, it is a ket with a very special relationship to \mathbf{M} . When $|\lambda\rangle$ is fed into the machine \mathbf{M} , all that happens is that it gets multiplied by the number λ . I'll give you an example. If \mathbf{M} is the 2×2 matrix

$$\begin{pmatrix} 1 & 2 \\ 2 & 1 \end{pmatrix},$$

then it's easy to see that the vector

$$\begin{pmatrix} 1 \\ 1 \end{pmatrix}$$

just gets multiplied by 3 when \mathbf{M} acts on it. Try it out. \mathbf{M} also happens to have another eigenvector:

$$\begin{pmatrix} 1 \\ -1 \end{pmatrix}.$$

When \mathbf{M} acts on this eigenvector, it multiplies the vector by a different number, namely -1 . On the other hand, if \mathbf{M} acts on the vector

$$\begin{pmatrix} 1 \\ 0 \end{pmatrix},$$

the vector is *not* simply multiplied by a number. \mathbf{M} alters the direction of the vector as well as its magnitude.

Just as the vectors that get multiplied by numbers when \mathbf{M} acts on them are called eigenvectors of \mathbf{M} , the constants that multiply them are called *eigenvalues*. In general, the eigenvalues are complex numbers. Here is an example that you can work out for yourself. Take the matrix

$$\mathbf{M} = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$$

and show that the vector

$$\begin{pmatrix} 1 \\ i \end{pmatrix}$$

is an eigenvector with eigenvalue $-i$.

Linear operators can also act on bra-vectors. The notation for multiplying $\langle B|$ by \mathbf{M} is

$$\langle B|\mathbf{M}.$$

I will keep the discussion short by telling you the rule for this type of multiplication. It is most simple in component form.

Remember that bra-vectors are represented in component form as row vectors. For example, the bra $\langle B|$ might be represented by

$$\langle B| = (\beta_1^* \quad \beta_2^* \quad \beta_3^*).$$

The rule is again just matrix multiplication. With a slight abuse of notation,

$$\langle B|\mathbf{M} = (\beta_1^* \quad \beta_2^* \quad \beta_3^*) \begin{pmatrix} m_{11} & m_{12} & m_{13} \\ m_{21} & m_{22} & m_{23} \\ m_{31} & m_{32} & m_{33} \end{pmatrix}. \quad (3.6)$$

3.1.3 Hermitian Conjugation

You might think that if $\mathbf{M}|A\rangle = |B\rangle$ then $\langle A|\mathbf{M} = \langle B|$, but if you do you are wrong. The problem is complex conjugation. Even when Z is just a complex number, if $Z|A\rangle = |B\rangle$, it is not generally true that $\langle A|Z = \langle B|$. You have to complex-conjugate Z when going from kets to bras: $\langle A|Z^* = \langle B|$. Of course, if Z happens to be a real number, then complex conjugation has no effect—every real number is equal to its own complex conjugate.

What we need is a concept of complex conjugation for operators. Let's look at the equation $\mathbf{M}|A\rangle = |B\rangle$ in component notation,

$$\sum_i m_{ji} \alpha_i = \beta_j,$$

and form its complex conjugate,

$$\sum_i m_{ji}^* \alpha_i^* = \beta_j^*.$$

We would like to write this equation in matrix form, using bras instead of kets. In doing this, we have to remember that bra-vectors are represented by rows, not columns. For the result to work out correctly, we also need to rearrange the complex conjugate elements of the matrix \mathbf{M} . The notation for this rearrangement is \mathbf{M}^\dagger , as explained below. Our new equation is

$$\langle A | \mathbf{M}^\dagger = \left(\alpha_1^* \quad \alpha_2^* \quad \alpha_3^* \right) \begin{pmatrix} m_{11}^* & m_{21}^* & m_{31}^* \\ m_{12}^* & m_{22}^* & m_{32}^* \\ m_{13}^* & m_{23}^* & m_{33}^* \end{pmatrix}. \quad (3.7)$$

Look carefully at the difference between the matrix in this equation and the matrix in Eq. 3.6. You will see two differences. The most obvious is the complex conjugation of each element, but you can also see a difference in the element indices. For example, where you see m_{23} in Eq. 3.6, you see m_{32}^* in Eq. 3.7. In other words, the rows and columns have been interchanged.

When we change an equation from the ket form to the bra form, we must modify the matrix in two steps:

1. Interchange the rows and the columns.
2. Complex-conjugate each matrix element.

In matrix notation, interchanging rows and columns is called *transposing* and is indicated by a superscript T . Thus, the transpose of the matrix \mathbf{M} is

$$\begin{pmatrix} m_{11} & m_{12} & m_{13} \\ m_{21} & m_{22} & m_{23} \\ m_{31} & m_{32} & m_{33} \end{pmatrix}^T = \begin{pmatrix} m_{11} & m_{21} & m_{31} \\ m_{12} & m_{22} & m_{32} \\ m_{13} & m_{23} & m_{33} \end{pmatrix}.$$

Notice that transposing a matrix flips it about the main diagonal (the diagonal from the upper left to the lower right).

The complex conjugate of a transposed matrix is called its *Hermitian conjugate*, denoted by a dagger. You could think of the dagger as a hybrid of the star-notation used in complex conjugation and the T used in transposition. In symbols,

$$\mathbf{M}^\dagger = [\mathbf{M}^T]^*.$$

To summarize: if \mathbf{M} acts on the ket $|A\rangle$ to give $|B\rangle$, then it follows that \mathbf{M}^\dagger acts on the bra $\langle A|$ to give $\langle B|$. In symbols:

If

$$\mathbf{M}|A\rangle = |B\rangle,$$

then

$$\langle A|\mathbf{M}^\dagger = \langle B|.$$

3.1.4 Hermitian Operators

Real numbers play a special role in physics. The results of any measurements are real numbers. Sometimes, we measure two quantities, put them together with an i (forming a

complex number), and call this number the result of a measurement. But it's actually just a way of combining two real measurements. If we want to be pedantic, we might say that observable quantities are equal to their own complex conjugates. That's of course just a fancy way of saying they are real. We are going to find out very soon that quantum mechanical observables are represented by linear operators. What kind of linear operators? The kind that are the closest thing to a real operator. Observables in quantum mechanics are represented by linear operators that are equal to their own Hermitian conjugates. They are called *Hermitian operators* after the French mathematician Charles Hermite. Hermitian operators satisfy the property

$$\mathbf{M} = \mathbf{M}^\dagger.$$

In terms of matrix elements, this can be stated as

$$m_{ji} = m_{ij}^*.$$

In other words, if you flip a Hermitian matrix about the main diagonal and then take its complex conjugate, the result is the same as the original matrix. Hermitian operators (and matrices) have some special properties. The first is that their eigenvalues are all real. Let's prove it.

Suppose λ and $|\lambda\rangle$ represent an eigenvalue and the corresponding eigenvector of the Hermitian operator \mathbf{L} . In symbols,

$$\mathbf{L}|\lambda\rangle = \lambda|\lambda\rangle.$$

Then, by the definition of Hermitian conjugation,

$$\langle\lambda|\mathbf{L}^\dagger = \langle\lambda|\lambda^*.$$

However, since \mathbf{L} is Hermitian, it is equal to \mathbf{L}^\dagger . Thus, we can rewrite the two equations as

$$\mathbf{L}|\lambda\rangle = \lambda|\lambda\rangle \tag{3.8}$$

and

$$\langle\lambda|\mathbf{L} = \langle\lambda|\lambda^*. \tag{3.9}$$

Now multiply Eq. 3.8 by $\langle\lambda|$ and Eq. 3.9 by $|\lambda\rangle$. They become

$$\langle\lambda|\mathbf{L}|\lambda\rangle = \lambda \langle\lambda|\lambda\rangle$$

and

$$\langle\lambda|\mathbf{L}|\lambda\rangle = \lambda^* \langle\lambda|\lambda\rangle.$$

Obviously, for both equations to be true, λ must equal λ^* . In other words, λ (and therefore any eigenvalue of a Hermitian operator) must be real.

3.1.5 Hermitian Operators and Orthonormal Bases

We come now to the basic mathematical theorem—I will call it the *fundamental theorem*—that serves as a foundation of quantum mechanics. The basic idea is that *observable quantities in quantum mechanics are represented by Hermitian operators*. It's a very simple theorem, but it's an extremely important one. We can state it more precisely as follows:

The Fundamental Theorem

- The eigenvectors of a Hermitian operator are a complete set. This means that **any vector the operator can generate** can be expanded as a sum of its eigenvectors.
- If λ_1 and λ_2 are two unequal eigenvalues of a Hermitian operator, then the corresponding eigenvectors are orthogonal.
- Even if the two eigenvalues are equal, the corresponding eigenvectors can be chosen to be orthogonal. This situation, where two different eigenvectors have the same eigenvalue, has a name: it's called *degeneracy*. Degeneracy comes into play when two operators have simultaneous eigenvectors, as discussed later on in Section 5.1.

One can summarize the fundamental theorem as follows: The eigenvectors of a Hermitian operator form an orthonormal basis. Let's prove it, beginning with the second bullet item.

According to the definition of eigenvectors and eigenvalues, we can write

$$\begin{aligned}\mathbf{L}|\lambda_1\rangle &= \lambda_1|\lambda_1\rangle \\ \mathbf{L}|\lambda_2\rangle &= \lambda_2|\lambda_2\rangle.\end{aligned}$$

Now, using the fact that \mathbf{L} is Hermitian (its own Hermitian conjugate), we can flip the first equation into a bra equation. Thus,

$$\begin{aligned}\langle\lambda_1|\mathbf{L} &= \lambda_1\langle\lambda_1| \\ \mathbf{L}|\lambda_2\rangle &= \lambda_2|\lambda_2\rangle.\end{aligned}$$

By now, the trick should be obvious, but I will spell it out. Take the first equation and form its inner product with $|\lambda_2\rangle$. Then, take the second equation and form its inner product with $\langle\lambda_1|$. The result is

$$\begin{aligned}\langle\lambda_1|\mathbf{L}|\lambda_2\rangle &= \lambda_1\langle\lambda_1|\lambda_2\rangle \\ \langle\lambda_1|\mathbf{L}|\lambda_2\rangle &= \lambda_2\langle\lambda_1|\lambda_2\rangle.\end{aligned}$$

By subtracting, we get

$$(\lambda_1 - \lambda_2)\langle\lambda_1|\lambda_2\rangle = 0.$$

Therefore, if λ_1 and λ_2 are different, the inner product $\langle\lambda_1|\lambda_2\rangle$ must be zero. In other words, the two eigenvectors must be orthogonal.

Next, let's prove that even if $\lambda_1 = \lambda_2$, the two eigenvectors can be *chosen* to be orthogonal. Suppose

$$\begin{aligned}\mathbf{L}|\lambda_1\rangle &= \lambda|\lambda_1\rangle \\ \mathbf{L}|\lambda_2\rangle &= \lambda|\lambda_2\rangle.\end{aligned}\tag{3.10}$$

In other words, there are two distinct eigenvectors with the same eigenvalue. It should be clear that any linear combination of the two eigenvectors is also an eigenvector with the same eigenvalue. With this much freedom, it is always possible to find two orthogonal linear combinations.

Let's see how. Consider an arbitrary linear combination of these two eigenvectors:

$$|A\rangle = \alpha|\lambda_1\rangle + \beta|\lambda_2\rangle.$$

Operating on both sides with \mathbf{L} , we get

$$\mathbf{L}|A\rangle = \alpha\mathbf{L}|\lambda_1\rangle + \beta\mathbf{L}|\lambda_2\rangle,$$

$$\mathbf{L}|A\rangle = \alpha\lambda|\lambda_1\rangle + \beta\lambda|\lambda_2\rangle,$$

and finally

$$\mathbf{L}|A\rangle = \lambda(\alpha|\lambda_1\rangle + \beta|\lambda_2\rangle) = \lambda|A\rangle.$$

This equation demonstrates that any linear combination of $|\lambda_1\rangle$ and $|\lambda_2\rangle$ is also an eigenvector of \mathbf{L} , with the same

eigenvalue. By assumption, these two vectors are linearly independent—otherwise, they would not represent distinct states. We will also suppose that they span the subspace of eigenvectors of \mathbf{L} that have eigenvalue λ . There is a straightforward process, called the *Gram-Schmidt* procedure, for finding an orthonormal basis for a subspace, given a set of independent vectors that spans the subspace. In plain English, we can find two orthonormal eigenvectors by writing them as a linear combination of $|\lambda_1\rangle$ and $|\lambda_2\rangle$. We outline the Gram-Schmidt procedure below, in Section 3.1.6.

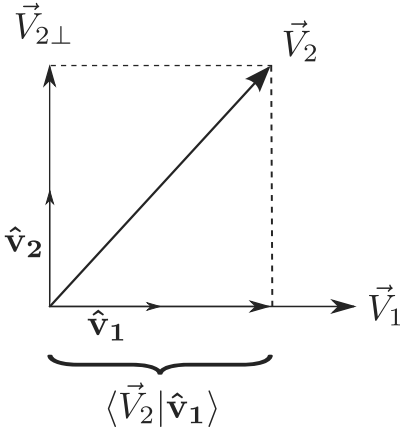
The final part of the theorem states that the eigenvectors are complete. In other words, if the space is N -dimensional, there will be N orthonormal eigenvectors. The proof is easy and I will leave it to you.

Exercise 3.1: Prove the following: If a vector space is N -dimensional, an orthonormal basis of N vectors can be constructed from the eigenvectors of a Hermitian operator.

3.1.6 The Gram-Schmidt Procedure

Sometimes we encounter a set of linearly independent eigenvectors that *do not* form an orthonormal set. This typically happens when a system has degenerate states—distinct states that have the same eigenvalue. In that situation, we can always use the linearly independent vectors we have, to create an orthonormal set that spans the same space. The method is the Gram-Schmidt procedure I alluded to earlier. Fig. 3.1 illustrates how it works for the simple case of two

linearly independent vectors. We start with the two vectors \vec{V}_1 and \vec{V}_2 , and from these we construct two orthonormal vectors, \hat{v}_1 and \hat{v}_2 .



$$\hat{v}_1 = \frac{\vec{V}_1}{|\vec{V}_1|}$$

$$\vec{V}_{2\perp} = \vec{V}_2 - \langle \vec{V}_2 | \hat{v}_1 \rangle \hat{v}_1$$

$$\hat{v}_2 = \frac{\vec{V}_{2\perp}}{|\vec{V}_{2\perp}|}$$

Figure 3.1: The Gram-Schmidt Procedure. Given two linearly independent vectors, \vec{V}_1 and \vec{V}_2 , that are not necessarily orthogonal, we can construct two orthonormal vectors, \hat{v}_1 and \hat{v}_2 . $\vec{V}_{2\perp}$ is an intermediate result used in the construction process. We can extend this procedure to larger sets of linearly independent vectors.

The first step is to divide \vec{V}_1 by its own length, $|\vec{V}_1|$, which gives us a unit vector parallel to \vec{V}_1 . We'll call that unit vector \hat{v}_1 , and \hat{v}_1 becomes the first vector in our orthonormal set. Next, we project \vec{V}_2 onto the direction of \hat{v}_1 by forming the inner product $\langle \vec{V}_2 | \hat{v}_1 \rangle$. Now, we subtract $\langle \vec{V}_2 | \hat{v}_1 \rangle \hat{v}_1$ from \vec{V}_2 . We'll call the result of this subtraction $\vec{V}_{2\perp}$. You can see in Fig. 3.1 that $\vec{V}_{2\perp}$ is orthogonal to \hat{v}_1 . Lastly, we divide $\vec{V}_{2\perp}$ by its own length to form the second member of our

orthonormal set, $\hat{\mathbf{v}}_2$. It should be clear that we can extend this procedure to larger sets of linearly independent vectors in more dimensions. For instance, if we had a third linearly independent vector, say \vec{V}_3 , pointing out of the page, we would subtract its projections onto each of the unit vectors $\hat{\mathbf{v}}_1$ and $\hat{\mathbf{v}}_2$, and then divide the result by its own length.¹

3.2 The Principles

We are now fully prepared to state the principles of quantum mechanics, so without further ado, let's do it.

The principles all involve the idea of an observable, and they presuppose the existence of an underlying complex vector space whose vectors represent system states. In this lecture, we present the four principles that do not involve the evolution of state-vectors with time. In Lecture 4, we will add a fifth principle that addresses the time development of system states.

An observable could also be called a measurable. It's a thing that you can measure with a suitable apparatus. Earlier, we spoke about measuring the components of a spin, σ_x , σ_y , and σ_z . These are examples of observables. We'll come back to them, but first let's look at the principles:

- Principle 1: The observable or measurable quantities of quantum mechanics are represented by linear operators **L**.

¹In this example, the term *out of the page* does not mean \vec{V}_3 is necessarily orthogonal to the plane of the page. The ability to use nonorthogonal vectors as a starting point is the main feature of the Gram-Schmidt Procedure.

I realize that this is the kind of hopelessly abstract statement that makes people give up on quantum mechanics and take up surfing instead. Don't worry—its meaning will become clear by the end of the lecture.

We'll soon see that \mathbf{L} must also be Hermitian. Some authors regard this as a postulate, or basic principle. We have chosen instead to derive it from the other principles. The end result is the same either way: the operators that represent observables are Hermitian.

- Principle 2: The possible results of a measurement are the eigenvalues of the operator that represents the observable. We'll call these eigenvalues λ_i . The state for which the result of a measurement is unambiguously λ_i is the corresponding eigenvector $|\lambda_i\rangle$. Don't unpack your surfboard just yet.

Here's another way to say it: if the system is in the eigenstate $|\lambda_i\rangle$, the result of a measurement is *guaranteed* to be λ_i .

- Principle 3: Unambiguously distinguishable states are represented by orthogonal vectors.
- Principle 4: If $|A\rangle$ is the state-vector of a system, and the observable \mathbf{L} is measured, the probability to observe value λ_i is

$$P(\lambda_i) = \langle A|\lambda_i\rangle\langle\lambda_i|A\rangle. \quad (3.11)$$

I'll remind you that the λ_i are the eigenvalues of \mathbf{L} , and $|\lambda_i\rangle$ are the corresponding eigenvectors.

These brief statements are hardly self-explanatory, and we'll need to flesh them out. For the moment, let's accept the first item, namely that every observable is identified with a linear operator. We can already begin to see that an operator is a way of packaging up states along with their eigenvalues, which are the possible results of measuring those states. These ideas should become clear as we move forward.

Let's recall some important points from our earlier discussion of spins. First of all, the result of a measurement is generally statistically uncertain. However, for any given observable, there are particular states for which the result is absolutely certain. For example, if the spin-measuring apparatus \mathcal{A} is oriented along the z axis, the state $|u\rangle$ always leads to the value $\sigma_z = +1$. Likewise, the state $|d\rangle$ never gives anything but $\sigma_z = -1$. Principle 1 gives us a new way to look at these facts. It implies that each observable (σ_x , σ_y , and σ_z) is identified with a specific linear operator in the two-dimensional space of states describing the spin.

When an observable is measured, the result is always a real number drawn from a set of possible results. For example, if the energy of an atom is measured, the result will be one of the established energy levels of the atom. For the familiar case of the spin, the possible values of any of the components are ± 1 . The apparatus never gives any other result. Principle 2 defines the relation between the operator representing an observable and the possible numerical outputs of the measurement. Namely, the result of a measurement is always one of the eigenvalues of the corresponding operator. Thus, each component of the spin operator must have two

eigenvalues equal to ± 1 .²

Principle 3 is the most interesting. At least I find it so. It speaks of *unambiguously distinct states*, a key idea that we have already encountered. Two states are physically distinct if there is a measurement that can tell them apart without ambiguity. For example, $|u\rangle$ and $|d\rangle$ can be distinguished by measuring σ_z . If you are handed a spin and told that it is either in the state $|u\rangle$ or the state $|d\rangle$, to find out which of the two states is the right one, all you have to do is align \mathcal{A} with the z axis and measure σ_z . There is no possibility of a mistake. The same is true for $|l\rangle$ and $|r\rangle$. You can distinguish them by measuring σ_x .

But suppose instead that you are told the spin is in one of the two states, $|u\rangle$ or $|r\rangle$ (*up* or *right*). There is nothing you can measure that will unambiguously tell you the spin's true state. Measuring σ_z won't do it. If you get $\sigma_z = +1$, it is possible that the initial state was $|r\rangle$ since there is a 50 percent probability of getting this answer in the state $|r\rangle$. For that reason, $|u\rangle$ and $|d\rangle$ are said to be physically distinguishable, but $|u\rangle$ and $|r\rangle$ are not. One might say that the inner product of two states is a measure of the inability to distinguish them with certainty. Sometimes this inner product is called the *overlap*. Principle 3 requires physically distinct states to be represented by orthogonal state-vectors, that is, vectors with no overlap. Thus, for spin states, $\langle u|d\rangle = 0$ but $\langle u|r\rangle = \frac{1}{\sqrt{2}}$.

²We have not yet explained what we mean by a “component” of the spin operator. We will do so shortly.

Finally, Principle 4 quantifies these ideas in a rule that expresses the probabilities for various outcomes of an experiment. If we assume that a system has been prepared in state $|A\rangle$, and subsequently the observable \mathbf{L} is measured, then the outcome will be one of the eigenvalues λ_i of the operator \mathbf{L} . But, in general, there is no way to tell for certain which of these values will be observed. There is only a probability—let us call it $P(\lambda_i)$ —that the outcome will be λ_i . Principle 4 tells us how to calculate that probability, and it is expressed in terms of the overlap of $|A\rangle$ and $|\lambda_i\rangle$. More precisely, the probability is the square of the magnitude of the overlap:

$$P(\lambda_i) = |\langle A|\lambda_i\rangle|^2$$

or, equivalently,

$$P(\lambda_i) = \langle A|\lambda_i\rangle\langle\lambda_i|A\rangle.$$

You might be wondering why the probability is not the overlap itself. Why the square of the overlap? Keep in mind that the inner product of two vectors is not always positive, or even real. Probabilities, on the other hand, are both positive and real. So it would not make sense to identify $P(\lambda_i)$ with $\langle A|\lambda_i\rangle$. But the square of the magnitude, $\langle A|\lambda_i\rangle\langle\lambda_i|A\rangle$, is always positive and real and thus can be identified with the probability of a given outcome.

An important consequence of the principles is as follows:

The operators that represent observables are Hermitian.

The reason for this is twofold. First, since the result of an experiment must be a real number, the eigenvalues of an operator \mathbf{L} must also be real. Secondly, the eigenvectors that represent unambiguously distinguishable results must have different eigenvalues, and must also be orthogonal. These conditions are sufficient to prove that \mathbf{L} must be Hermitian.

3.3 An Example: Spin Operators

It may be hard to believe, but single spins—as simple as they are—still have a lot more to teach us about quantum mechanics, and we plan to milk them for all they’re worth. Our goal in this section is to write down the spin operators in concrete form, as 2×2 matrices. Then, we’ll get to see how they work in specific situations. We’ll build up our spin operators and state-vectors shortly. But before we dive into the details, I’d like to say a little more about how operators are related to physical measurements. The relationship is a subtle one, and we’ll say more about it as we go.

As you know, physicists recognize various types of physical quantities, such as scalars and vectors. It should come as no surprise, then, that an operator associated with the measurement of a vector (such as spin) has a vector character of its own.

In our travels so far, we have seen more than one kind of vector. The 3-vector is the most straightforward and serves as a prototype. It’s a mathematical representation of an arrow in three-dimensional space, and is often represented by three real numbers, written out as a column matrix. Because their components are real-valued, 3-vectors are not quite rich

enough to represent quantum states. For that, we need bras and kets, which have complex-valued components.

What sort of vector is the spin operator σ ? It is definitely *not* a state-vector (a bra or a ket). It's not exactly a 3-vector either, but it does have a strong family resemblance because it's associated with a direction in space. In fact, we will frequently use σ as though it were a simple 3-vector. However, we'll try to keep things straight by calling σ a *3-vector operator*.

But what does that actually mean? In physical terms, it means this: Just as a spin-measuring apparatus can only *answer questions* about a spin's orientation in a specific direction, a spin operator can only provide *information* about the spin component in a specific direction. To physically measure spin in a different direction, we need to rotate the apparatus to point in the new direction. The same idea applies to the spin operator—if we want it to tell us about the spin component in a new direction, it too must be “rotated,” but this kind of rotation is accomplished mathematically. The bottom line is that there is a spin operator for each direction in which the apparatus can be oriented.

3.4 Constructing Spin Operators

Now, let's work out the details of spin operators. The first goal is to construct operators to represent the components of spin, σ_x , σ_y , and σ_z . Then we'll build on those results to construct an operator that represents a spin component in any direction. As usual, we begin with σ_z . We know that σ_z has definite, unambiguous values for the states $|u\rangle$ and $|d\rangle$,

and that the corresponding measurement values are $\sigma_z = +1$ and $\sigma_z = -1$. Here is what the first three principles tell us:

- Principle 1: Each component of σ is represented by a linear operator.
- Principle 2: The eigenvectors of σ_z are $|u\rangle$ and $|d\rangle$. The corresponding eigenvalues are $+1$ and -1 . We can express this with the abstract equations

$$\begin{aligned}\sigma_z|u\rangle &= |u\rangle \\ \sigma_z|d\rangle &= -|d\rangle.\end{aligned}\tag{3.12}$$

- Principle 3: States $|u\rangle$ and $|d\rangle$ are orthogonal to each other. This can be expressed as

$$\langle u|d\rangle = 0.\tag{3.13}$$

Recalling our column representations of $|u\rangle$ and $|d\rangle$ from Eqs. 2.11 and 2.12, we can write Eqs. 3.12 in matrix form as

$$\begin{pmatrix} (\sigma_z)_{11} & (\sigma_z)_{12} \\ (\sigma_z)_{21} & (\sigma_z)_{22} \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}\tag{3.14}$$

and

$$\begin{pmatrix} (\sigma_z)_{11} & (\sigma_z)_{12} \\ (\sigma_z)_{21} & (\sigma_z)_{22} \end{pmatrix} \begin{pmatrix} 0 \\ 1 \end{pmatrix} = -\begin{pmatrix} 0 \\ 1 \end{pmatrix}.\tag{3.15}$$

There is only one matrix that satisfies these equations. I leave it as an exercise to prove

$$\begin{pmatrix} (\sigma_z)_{11} & (\sigma_z)_{12} \\ (\sigma_z)_{21} & (\sigma_z)_{22} \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \quad (3.16)$$

or, more concisely,

$$\sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \quad (3.17)$$

Exercise 3.2: Prove that Eq. 3.16 is the unique solution to Eqs. 3.14 and 3.15.

This is our very first example of a quantum mechanical operator. Let's summarize what went into it. First, some experimental data: there are certain states that we called $|u\rangle$ and $|d\rangle$, in which the measurement of σ_z gives unambiguous results ± 1 . Next, the principles told us that $|u\rangle$ and $|d\rangle$ are orthogonal and are eigenvectors of a linear operator σ_z . Finally, we learned from the principles that the corresponding eigenvalues are the observed (or measured) values, again ± 1 . That's all it takes to derive Eq. 3.17.

Can we do the same for the other two components of spin, σ_x and σ_y ? Yes, we can.³ The eigenvectors of σ_x are $|r\rangle$ and $|l\rangle$, with eigenvalues $+1$ and -1 respectively. In equation form,

³We are not trying to slip in a political slogan. Really. Just say *no* to slogans.

$$\begin{aligned}\sigma_x|r\rangle &= |r\rangle \\ \sigma_x|l\rangle &= -|l\rangle.\end{aligned}\tag{3.18}$$

Recall that $|r\rangle$ and $|l\rangle$ are linear superpositions of $|u\rangle$ and $|d\rangle$:

$$\begin{aligned}|r\rangle &= \frac{1}{\sqrt{2}}|u\rangle + \frac{1}{\sqrt{2}}|d\rangle \\ |l\rangle &= \frac{1}{\sqrt{2}}|u\rangle - \frac{1}{\sqrt{2}}|d\rangle.\end{aligned}\tag{3.19}$$

Substituting the appropriate column vectors for $|u\rangle$ and $|d\rangle$, we get

$$\begin{aligned}|r\rangle &= \begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix} \\ |l\rangle &= \begin{pmatrix} \frac{1}{\sqrt{2}} \\ -\frac{1}{\sqrt{2}} \end{pmatrix}.\end{aligned}$$

To make Eqs. 3.18 concrete, we can write them in matrix form:

$$\begin{pmatrix} (\sigma_x)_{11} & (\sigma_x)_{12} \\ (\sigma_x)_{21} & (\sigma_x)_{22} \end{pmatrix} \begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix} = \begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix}$$

and

$$\begin{pmatrix} (\sigma_x)_{11} & (\sigma_x)_{12} \\ (\sigma_x)_{21} & (\sigma_x)_{22} \end{pmatrix} \begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{-1}{\sqrt{2}} \end{pmatrix} = - \begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{-1}{\sqrt{2}} \end{pmatrix}.$$

If you write these equations out in longhand form, they turn into four easily solved equations for the matrix elements $(\sigma_x)_{11}$, $(\sigma_x)_{12}$, $(\sigma_x)_{21}$, and $(\sigma_x)_{22}$. Here is the solution:

$$\begin{pmatrix} (\sigma_x)_{11} & (\sigma_x)_{12} \\ (\sigma_x)_{21} & (\sigma_x)_{22} \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

or

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

Finally, we can do the same for σ_y . The eigenvectors of σ_y are the *in* and *out* states $|i\rangle$ and $|o\rangle$:

$$\begin{aligned} |i\rangle &= \frac{1}{\sqrt{2}}|u\rangle + \frac{i}{\sqrt{2}}|d\rangle \\ |o\rangle &= \frac{1}{\sqrt{2}}|u\rangle - \frac{i}{\sqrt{2}}|d\rangle. \end{aligned}$$

In component form, these equations become

$$\begin{aligned} |i\rangle &= \begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{i}{\sqrt{2}} \end{pmatrix} \\ |o\rangle &= \begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{-i}{\sqrt{2}} \end{pmatrix}, \end{aligned}$$

and an easy calculation gives

$$\sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}.$$

To summarize, the three operators σ_x , σ_y , and σ_z are represented by the three matrices

$$\begin{aligned} \sigma_z &= \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \\ \sigma_x &= \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \\ \sigma_y &= \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}. \end{aligned} \tag{3.20}$$

These three matrices are very famous and carry the name of their discoverer. They are the *Pauli matrices*.⁴

3.5 A Common Misconception

This is a convenient time to warn you about a potential hazard. The correspondence between operators and measurements is fundamental in quantum mechanics. It is also very easy to misunderstand. Here's what is true about operators in quantum mechanics:

1. Operators are the things we use to calculate eigenvalues and eigenvectors.
2. Operators act on state-vectors (which are abstract mathematical objects), not on actual physical systems.

⁴Along with the 2×2 identity matrix, they are also *quaternions*.

3. When an operator acts on a state-vector, it produces a new state-vector.

Having said what is true about operators, I want to warn you about a common misconception. It is often thought that measuring an observable is the same as operating with the corresponding operator on the state. For example, suppose we are interested in measuring an observable \mathbf{L} . The measurement is some kind of operation that the apparatus does to the system, but that operation is in no way the same as acting on the state with the operator \mathbf{L} . For example, if the state of the system before we do the measurement is $|A\rangle$, it is *not* correct to say that the measurement of \mathbf{L} changes the state to $\mathbf{L}|A\rangle$.

To make sense of this, let's look closely at an example. Fortunately, the spin example of the previous subsection is just what we need. Recall Eqs. 3.12:

$$\sigma_z|u\rangle = |u\rangle$$

$$\sigma_z|d\rangle = -|d\rangle.$$

In these situations, there is no trap because $|u\rangle$ and $|d\rangle$ are eigenvectors of σ_z . If the system is prepared in, say, the $|d\rangle$ state, a measurement will definitely give the result -1 , and the σ_z operator transforms the prepared state into the corresponding post-measurement state, $-|d\rangle$. The state $-|d\rangle$ is the same as $|d\rangle$ except for a multiplicative constant, so the two states are really the same. No problems here.

But now let's review the action of σ_z on the prepared state $|r\rangle$, which is *not* one of its eigenvectors. From Eq. 3.19, we know that

$$|r\rangle = \frac{1}{\sqrt{2}}|u\rangle + \frac{1}{\sqrt{2}}|d\rangle.$$

Acting on this state-vector with σ_z gives the result

$$\sigma_z|r\rangle = \frac{1}{\sqrt{2}}\sigma_z|u\rangle + \frac{1}{\sqrt{2}}\sigma_z|d\rangle$$

or

$$\sigma_z|r\rangle = \frac{1}{\sqrt{2}}|u\rangle - \frac{1}{\sqrt{2}}|d\rangle. \quad (3.21)$$

OK, here is our trap. Despite what you might think, the state-vector on the right-hand side of Eq. 3.21 is definitely *not* the state that would result from a measurement of σ_z . That measurement result would be either $+1$, leaving the system in state $|u\rangle$, or -1 , leaving it in state $|d\rangle$. Neither of these results would leave the system state-vector in the superposition represented by Eq. 3.21.

But surely that state-vector must have *something* to do with the measurement result? In fact, it does. We'll find part of the answer in Lecture 4, where we'll see how the new state-vector allows us to calculate the probabilities of each possible outcome of the measurement. However, the result of a measurement cannot be properly described without taking the apparatus into account as part of the system. What actually does happen during a measurement is the subject of Section 7.8.

3.6 3-Vector Operators Revisited

Now, let's revisit the idea of a 3-vector operator. I have called σ_x , σ_y , and σ_z the components of spin along the three axes, implying that they are the components of some kind of 3-vector. This is a good time to return to the two notions of vectors that come up all the time in physics. First, there is your garden-variety vector in ordinary three-dimensional space, which we've decided to call a 3-vector. As we've seen, a 3-vector has components along the three directions of space.

The other completely distinct meaning of the term *vector* is the state-vector of a system. Thus, $|u\rangle$ and $|d\rangle$, $|r\rangle$ and $|l\rangle$, and $|i\rangle$ and $|o\rangle$ are state-vectors in a two-dimensional space of spin states. What about σ_x , σ_y , and σ_z ? Are they vectors, and if so, what kind?

Clearly, they are not state-vectors; they are operators (written as matrices) that correspond to the three measurable components of spin. In fact, these 3-vector operators represent a new type of vector. They are different both from state-vectors, and from ordinary 3-vectors. However, because spin operators behave so much like 3-vectors, it does no harm to think of them in that way, and that's what we'll do here.

We measure spin components by orienting the apparatus \mathcal{A} along any one of the three axes and then activating it. But then why not orient \mathcal{A} along *any* axis and measure the component of σ along that axis? In other words, take any unit 3-vector \hat{n} with components n_x , n_y , and n_z , and orient

the apparatus \mathcal{A} with its arrow along \hat{n} . Activating \mathcal{A} would then measure the component of σ along the axis \hat{n} . There must be an operator that corresponds to this measurable quantity.

If σ really behaves like a 3-vector, then the component of σ along \hat{n} is nothing but the ordinary dot product of σ and \hat{n} .^{5,6} Let's denote that component of σ by σ_n , so that

$$\sigma_n = \vec{\sigma} \cdot \hat{n}$$

or, in expanded form,

$$\sigma_n = \sigma_x n_x + \sigma_y n_y + \sigma_z n_z. \quad (3.22)$$

To clarify the meaning of this equation, keep in mind that the components of \hat{n} are just numbers. They themselves are not operators. Eq. 3.22 describes a vector-operator that is constructed as the sum of three terms, each containing a numerical coefficient n_x , n_y , or n_z . To be more concrete, we can write Eq. 3.22 in matrix form:

$$\sigma_n = n_x \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} + n_y \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} + n_z \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

⁵We'll start using the notation $\vec{\sigma}$, except when referring to components, such as σ_x .

⁶The careful reader may object, because the result of this "ordinary" dot product is a 2×2 matrix rather than a scalar, so it's not *quite* ordinary. Perhaps there is some comfort in the fact that the resulting matrix operator corresponds to a vector component, which *is* a scalar. It all works out in the end.

Or even more explicitly, we can combine these three terms into a single matrix:

$$\sigma_n = \begin{pmatrix} n_z & (n_x - in_y) \\ (n_x + in_y) & -n_z \end{pmatrix}. \quad (3.23)$$

What is this good for? Not much, until we find the eigenvectors and eigenvalues of σ_n . But once we do that, we will know the possible outcomes of a measurement along the direction of \hat{n} . And we will also be able to calculate probabilities for those outcomes. In other words, we will have a complete picture of spin measurements in three-dimensional space. That is pretty darn cool, if I say so myself.

3.7 Reaping the Results

We are now positioned to make some real calculations, something that should make your inner physicist jump for joy. Let's look at the special case where \hat{n} lies in the x - z plane, which is the plane of this page. Since \hat{n} is a unit vector, we can write

$$n_z = \cos \theta$$

$$n_x = \sin \theta$$

$$n_y = 0,$$

where θ is the angle between the z axis and the \hat{n} axis. Plugging these values into Eq. 3.23, we can write

$$\sigma_n = \begin{pmatrix} \cos \theta & \sin \theta \\ \sin \theta & -\cos \theta \end{pmatrix}.$$

Exercise 3.3: Calculate the eigenvectors and eigenvalues of σ_n . *Hint:* Assume the eigenvector λ_1 has the form

$$\begin{pmatrix} \cos \alpha \\ \sin \alpha \end{pmatrix},$$

where α is an unknown parameter. Plug this vector into the eigenvalue equation and solve for α in terms of θ . Why did we use a single parameter α ? Notice that our suggested column vector must have unit length.

Here are the results:

$$\lambda_1 = 1$$

$$|\lambda_1\rangle = \begin{pmatrix} \cos \frac{\theta}{2} \\ \sin \frac{\theta}{2} \end{pmatrix}$$

and

$$\lambda_2 = -1$$

$$|\lambda_2\rangle = \begin{pmatrix} -\sin \frac{\theta}{2} \\ \cos \frac{\theta}{2} \end{pmatrix}.$$

Notice some important facts. First, the two eigenvalues are again $+1$ and -1 . This should come as no surprise; the apparatus \mathcal{A} can only give one of these two answers no matter which way it points. But it's good to see this come out of the equations. The second fact is that the two eigenvectors are orthogonal.

We are now ready to make an experimental prediction. Suppose \mathcal{A} initially points along the z axis and that we prepare a spin in the up state $|u\rangle$. Then, we rotate \mathcal{A} so that it lies along the \hat{n} axis. What is the probability of observing $\sigma_n = +1$? According to Principle 4, and using the row and column expansions of $\langle u|$ and $|\lambda_1\rangle$, the answer is

$$P(+1) = |\langle u|\lambda_1\rangle|^2 = \cos^2 \frac{\theta}{2}. \quad (3.24)$$

Similarly, for the same setup,

$$P(-1) = |\langle u|\lambda_2\rangle|^2 = \sin^2 \frac{\theta}{2}. \quad (3.25)$$

With this result, we have come nearly full circle. When introducing spins, we made the claim that if we prepare a large number of them in the up state and then measure their component along \hat{n} , at angle θ to the z axis, then the average value of the measured results would be $\cos \theta$ —the same result we would get for a simple 3-vector in classical physics. Does our mathematical framework give the same result? *It had better!* If a theory disagrees with experiment, it's the theory that has to leave town. Let's see how well our theory holds up so far.

Unfortunately, we need to cheat a little by using an equation that we will not fully explain until the next lecture. This is the equation that tells us how to calculate the average value (also called the expectation value) of a measurement. Here it is:

$$\langle \mathbf{L} \rangle = \sum_i \lambda_i P(\lambda_i). \quad (3.26)$$

It's worth mentioning that Eq. 3.26 is just a standard formula for an average value. It's not unique to quantum mechanics.

To calculate the expectation value of a measurement corresponding to the operator \mathbf{L} , we multiply each eigenvalue by its probability, and then sum the results. Of course, the operator we're looking at now is just σ_n , and we already have all the values we need. Let's plug them in. Using Eqs. 3.24 and 3.25, along with our known eigenvalues, we can write

$$\langle \sigma_n \rangle = (+1) \cos^2 \frac{\theta}{2} + (-1) \sin^2 \frac{\theta}{2}$$

or

$$\langle \sigma_n \rangle = \cos^2 \frac{\theta}{2} - \sin^2 \frac{\theta}{2}.$$

If you remember your trigonometry, this gives

$$\langle \sigma_n \rangle = \cos \theta,$$

which agrees perfectly with experiment. Yes! We've done it!

Having come this far, you might want to try your hand on a slightly more general problem. As before, we start with the apparatus \mathcal{A} pointing in the z direction. But now, once the spin has been prepared in the up state, we can rotate \mathcal{A} to an arbitrary direction in space for the second set of measurements. In this situation, $n_y \neq 0$. Go ahead and try it.

Exercise 3.4: Let $n_z = \cos \theta$, $n_x = \sin \theta \cos \phi$, and $n_y = \sin \theta \sin \phi$. Angles θ and ϕ are defined according to the usual conventions for spherical coordinates (Fig. 3.2). Compute the eigenvalues and eigenvectors for the matrix of Eq. 3.23.

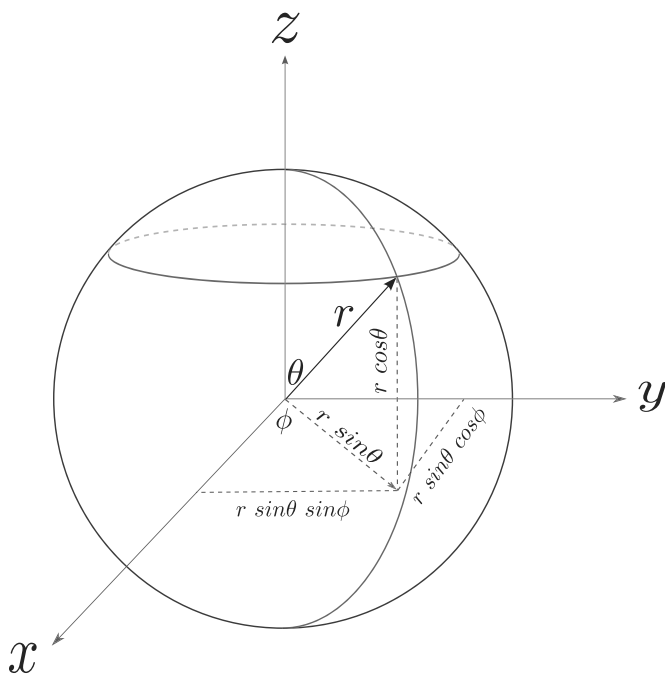


Figure 3.2: Spherical Coordinates. This diagram illustrates conventional spherical coordinate labels r , θ , and ϕ . It also illustrates the conversion to Cartesian coordinates: $x = r \sin \theta \cos \phi$, $y = r \sin \theta \sin \phi$, and $z = r \cos \theta$.

You could also try working out a much more elaborate example involving two directions, \hat{n} and \hat{m} . In this setup, \mathcal{A} not only *ends up* in an arbitrary direction; it also *starts out* in a (different) arbitrary direction.

Exercise 3.5: Suppose that a spin is prepared so that $\sigma_m = +1$. The apparatus is then rotated to the \hat{n} direction and σ_n is measured. What is the probability that the result is $+1$? Note that $\sigma_m = \sigma \cdot \hat{m}$, using the same convention we used for σ_n .

The answer is the square of the cosine of half the angle between \hat{m} and \hat{n} . Can you show it?

3.8 The Spin-Polarization Principle

There is an important theorem that you can try to prove. I will call it

The Spin-Polarization Principle: *Any state of a single spin is an eigenvector of some component of the spin.*

In other words, given any state

$$|A\rangle = \alpha_u|u\rangle + \alpha_d|d\rangle,$$

there exists some direction \hat{n} , such that

$$\vec{\sigma} \cdot \vec{n} |A\rangle = |A\rangle.$$

This means that for any spin state, there is some orientation of the apparatus \mathcal{A} such that \mathcal{A} will register $+1$ when it acts. In physics language, we say that the states of a spin are characterized by a *polarization vector*, and along that polarization vector the component of spin is predictably $+1$, assuming of course that you know the state-vector.

An interesting consequence of this theorem is that there is no state for which the *expectation values* of all three components of spin are zero. There is a quantitative way to express this. Consider the expectation value of the spin along the direction \hat{n} . Since $|A\rangle$ is an eigenvector of $\vec{\sigma} \cdot \vec{n}$ (with eigenvalue $+1$), it follows that the expectation value can be expressed as

$$\langle \vec{\sigma} \cdot \vec{n} \rangle = 1.$$

On the other hand, the expectation value of the perpendicular components of σ are zero in the state $|A\rangle$. It follows that the squares of the expectation values of all three components of σ sum to 1. Moreover, this is true for any state:

$$\langle \sigma_x \rangle^2 + \langle \sigma_y \rangle^2 + \langle \sigma_z \rangle^2 = 1. \quad (3.27)$$

Remember this fact. We will come back to it in Lecture 6.

Lecture 4

Time and Change

There is a massive, quiet, intimidating man sitting alone at the end of the bar. His T-shirt says “-1.”

Art: *Who is that “Minus One” guy over in the corner? The bouncer?*

Lenny: *He’s way more than a bouncer. He’s*

THE LAW.

Without him, this whole place would fall apart.

4.1 A Classical Reminder

In *Volume I*, it took a little more than a page to explain what a state is in classical mechanics. The quantum version has taken three lectures, three mathematical interludes, and according to my rough count, about 17,000 words to get to the same place. But I think the worst is over. We now know what a state is. However, just as in classical physics,

knowing the states of a system is only half the story. The other half involves a rule about how states change with time. That's our next job.

Let me just give you a quick reminder about the nature of change in classical physics. In classical physics, the space of states is a mathematical set. The logic is Boolean, and the evolution of states over time is deterministic and reversible. In the simplest examples we considered, the state-space consisted of a few points: Heads and Tails for a coin, $\{1, 2, 3, 4, 5, 6\}$ for a die. The states were pictured as a set of points on the page, and the time evolution was just a rule telling you where to go next. A law of motion consisted of a graph with arrows connecting the states. The main rule—determinism—was that wherever you are in the state-space, the next state is completely specified by the law of motion. But there was also another rule called reversibility. Reversibility is the requirement that a properly formulated law must also tell you where you were last. A good law corresponds to a graph with exactly one arrow in and one arrow out at each state.

There is another way to describe these requirements. I called it the *minus first law*, because it underlies everything else. It says that information is never lost. If two identical isolated systems start out in different states, they stay in different states. Moreover, in the past they were also in different states. On the other hand, if two identical systems are in the same state at some point in time, then their histories and their future evolutions must also be identical. Distinctions are conserved. The quantum version of the minus first law has a name—*unitarity*.

4.2 Unitarity

Let us consider a closed system that at time t is in the quantum state $|\Psi\rangle$. (The use of the Greek letter Ψ [psi] for quantum states is traditional when considering the evolution of systems.) To indicate that the state was $|\Psi\rangle$ at the specific time t , let's complicate the notation a bit and call the state $|\Psi(t)\rangle$. Of course, this notation suggests a bit more than just "the state was $|\Psi\rangle$ at time t ." It also suggests that the state may be different at different times. Thus, we think of $|\Psi(t)\rangle$ as representing the entire history of the system.

The basic dynamical assumption of quantum mechanics is that if you know the state at one time, then the quantum equations of motion tell you what it will be later. Without loss of generality, we can take the initial time to be zero and the later time to be t . The state at time t is given by some operation that we call $\mathbf{U}(t)$, acting on the state at time zero. Without further specifying the properties of $\mathbf{U}(t)$, this tells us very little except that $|\Psi(t)\rangle$ is determined by $|\Psi(0)\rangle$. Let's express this relation with the equation

$$|\Psi(t)\rangle = \mathbf{U}(t)|\Psi(0)\rangle. \quad (4.1)$$

The operation \mathbf{U} is called *the time-development operator* for the system.

4.3 Determinism in Quantum Mechanics

At this point, we need to draw some careful distinctions. We are setting up $\mathbf{U}(t)$ in such a way that the state-vector will evolve in a deterministic manner. Yes, you heard me correctly—the time evolution of the state-vector is *deterministic*. This is nice because it provides us with something we can try to predict. But how does that square with the statistical character of our measurement results?

As we've seen, knowing the quantum state does not mean that you can predict the result of an experiment with certainty. For example, knowing that the state of a spin is $|r\rangle$ may tell you the outcome of a measurement of σ_x but tells you nothing about a measurement of σ_z or σ_y . For this reason, Eq. 4.1 is not the same as classical determinism. Classical determinism allows us to predict the results of experiments. The quantum evolution of states allows us to compute the probabilities of the outcomes of *later* experiments.

This is one of the core differences between classical and quantum mechanics. It goes back to the relationship between states and measurements we mentioned at the very beginning of this book. In classical mechanics, there's no real difference between states and measurements. In quantum mechanics, the difference is profound.

4.4 A Closer Look at $\mathbf{U}(t)$

Conventional quantum mechanics places a couple of requirements on $\mathbf{U}(t)$. First, it requires $\mathbf{U}(t)$ to be a linear operator. That is not very surprising. The relationships between states in quantum mechanics are always linear. It goes along with the idea that the state-space is a vector space. But linearity is not the only thing that quantum mechanics requires of $\mathbf{U}(t)$. It also requires the quantum analog of the minus first law: *the conservation of distinctions*.

Recall from the last lecture that two states are distinguishable if they are orthogonal. Being orthogonal, two different basis vectors represent two distinguishable states. Suppose that $|\Psi(0)\rangle$ and $|\Phi(0)\rangle$ are two distinguishable states; in other words, there is a precise experiment that can tell them apart, and therefore they must be orthogonal:

$$\langle\Psi(0)|\Phi(0)\rangle = 0.$$

The conservation of distinctions implies that they will continue to be orthogonal for all time. We can express this as

$$\langle\Psi(t)|\Phi(t)\rangle = 0 \tag{4.2}$$

for all values of t . This principle has consequences for the time-development operator $\mathbf{U}(t)$. To see what they are, let's flip the ket-vector Eq. 4.1 to its bra-vector counterpart:

$$\langle\Psi(t)| = \langle\Psi(0)|\mathbf{U}^\dagger(t). \tag{4.3}$$

Notice the dagger that indicates Hermitian conjugation. Now, let's plug Eqs. 4.1 and 4.3 into Eq. 4.2:

$$\langle \Psi(0) | \mathbf{U}^\dagger(t) \mathbf{U}(t) | \Phi(0) \rangle = 0. \quad (4.4)$$

To examine the consequences of this equation, consider an orthonormal basis of vectors $|i\rangle$. Any basis will do. The orthonormality is expressed in equation form as

$$\langle i | j \rangle = \delta_{ij},$$

where δ_{ij} is the usual Kronecker symbol.

Next, let's take $|\Phi(0)\rangle$ and $|\Psi(0)\rangle$ to be members of this orthonormal basis. Substituting into Eq. 4.4 gives

$$\langle i | \mathbf{U}^\dagger(t) \mathbf{U}(t) | j \rangle = 0 \quad (i \neq j)$$

whenever i and j are not the same. On the other hand, if i and j are the same, then so are the output vectors $\mathbf{U}(t)|i\rangle$ and $\mathbf{U}(t)|j\rangle$. In that case, the inner product between them should be 1. Therefore, the general relation takes the form

$$\langle i | \mathbf{U}^\dagger(t) \mathbf{U}(t) | j \rangle = \delta_{ij}.$$

In other words, the operator $\mathbf{U}^\dagger(t) \mathbf{U}(t)$ behaves like the unit operator I when it acts between any members of a basis set. From here it is an easy matter to prove that $\mathbf{U}^\dagger(t) \mathbf{U}(t)$ acts like the unit operator I when it acts on any state. An operator \mathbf{U} that satisfies

$$\mathbf{U}^\dagger \mathbf{U} = I$$

is called *unitary*. In physics lingo, *time evolution is unitary*.

Unitary operators play an enormous role in quantum mechanics, representing all sorts of transformations on the state-space. Time evolution is just one example. Thus, we conclude this section with a fifth principle of quantum mechanics:

- Principle 5: The evolution of state-vectors with time is unitary.

Exercise 4.1: Prove that if \mathbf{U} is unitary, and if $|A\rangle$ and $|B\rangle$ are any two state-vectors, then the inner product of $\mathbf{U}|A\rangle$ and $\mathbf{U}|B\rangle$ is the same as the inner product of $|A\rangle$ and $|B\rangle$. One could call this the *conservation of overlaps*. It expresses the fact that the logical relation between states is preserved with time.

4.5 The Hamiltonian

In the study of classical mechanics, we became familiar with the idea of an incremental change in time. Quantum mechanics is no different in this respect: we may build up finite time intervals by combining many infinitesimal intervals. Doing so will lead to a differential equation for the evolution of the state-vector. To that end, we replace the time interval t with an infinitesimal time interval ϵ and consider the time-evolution operator for this small interval.

There are two principles that go into the study of incremental changes. The first principle is unitarity:

$$\mathbf{U}^\dagger(\epsilon)\mathbf{U}(\epsilon) = I. \quad (4.5)$$

The second principle is continuity. This means that the state-vector changes smoothly. To make this precise, first consider the case in which ϵ is zero. It should be obvious that in this case the time-evolution operator is merely the unit operator I . Continuity means that when ϵ is very small, $\mathbf{U}(\epsilon)$ is close to the unit operator, differing from it by something of order ϵ . Thus, we write

$$\mathbf{U}(\epsilon) = I - i\epsilon\mathbf{H}. \quad (4.6)$$

You may wonder why I put a minus sign and an i in front of \mathbf{H} . These factors are completely arbitrary at this stage. In other words, they are a convention that has no content. I used them with an eye toward the future, when we will recognize \mathbf{H} as something familiar from classical physics.

We will also need an expression for \mathbf{U}^\dagger . Remembering that Hermitian conjugation requires the complex conjugation of coefficients, we find that

$$\mathbf{U}^\dagger(\epsilon) = I + i\epsilon\mathbf{H}^\dagger. \quad (4.7)$$

Now we plug Eqs. 4.6 and 4.7 into the unitarity condition of Eq. 4.5:

$$(I + i\epsilon\mathbf{H}^\dagger)(I - i\epsilon\mathbf{H}) = I.$$

Expanding to first order in ϵ , we find

$$\mathbf{H}^\dagger - \mathbf{H} = 0$$

or, in a format that is more illuminating,

$$\mathbf{H}^\dagger = \mathbf{H}. \quad (4.8)$$

This last equation expresses the unitarity condition. But it also says that \mathbf{H} is a Hermitian operator. This has great significance. We can now say that \mathbf{H} is an observable, and has a complete set of orthonormal eigenvectors and eigenvalues. As we proceed, \mathbf{H} will become a very familiar object, namely the *quantum Hamiltonian*. Its eigenvalues are the values that would result from measuring the energy of a quantum system. Exactly why we identify \mathbf{H} with the classical concept of a Hamiltonian, and its eigenvalues with energy, will become clear shortly.

Let's return now to Eq. 4.1 and specialize it to the infinitesimal case $t = \epsilon$. Using Eq. 4.6, we find

$$|\Psi(\epsilon)\rangle = |\Psi(0)\rangle - i\epsilon\mathbf{H}|\Psi(0)\rangle.$$

This is just the kind of equation that we can easily turn into a differential equation. First, we transpose the first term on the right side over to the left side, and then divide by ϵ :

$$\frac{|\Psi(\epsilon)\rangle - |\Psi(0)\rangle}{\epsilon} = -i\mathbf{H}|\Psi(0)\rangle.$$

If you remember your calculus (see *Volume I* for a quick review), you'll recognize that the left-hand side of this equation looks exactly like the definition of a derivative. If we take the limit as $\epsilon \rightarrow 0$, it becomes the time derivative of the state-vector:

$$\frac{\partial|\Psi\rangle}{\partial t} = -i\mathbf{H}|\Psi\rangle. \quad (4.9)$$

We originally set things up so that the time variable was zero, but there was nothing special about $t = 0$. Had we chosen another time and done the same thing, we would have gotten exactly the same result, namely, Eq. 4.9. This equation tells us how the state-vector changes: if we know the state-vector at one instant, the equation tells us what it will be at the next. Eq. 4.9 is important enough to have a name. It is called the *generalized Schrödinger equation*, or more commonly, the *time-dependent Schrödinger equation*. If we know the Hamiltonian, it tells us how the state of an undisturbed system evolves with time. Art likes to call this state-vector *Schrödinger's Ket*. He even wanted to render the Greek symbol with little whiskers,¹ but I had to draw the line somewhere.

4.6 What Ever Happened to \hbar ?

I'm sure you have all heard of Planck's constant. Planck himself called it h and gave it a value of about $6.6 \times 10^{-34} \text{ kg m}^2 / \text{s}$.

¹OK, not really.

Later generations redefined it, dividing by a factor of 2π and calling the result \hbar :

$$\hbar = \frac{h}{2\pi} = 1.054571726 \dots \times 10^{-34} \text{ kg } m^2/s.$$

Why divide by 2π ? Because it saves us from having to write 2π in lots of other places. Considering the importance of Planck's constant in quantum mechanics, it seems a little odd that it hasn't come up yet. We're going to correct that now.

In quantum mechanics, as in classical physics, the Hamiltonian is the mathematical object that represents the energy of a system. This raises a question that, if you are very alert, may have been a source of confusion. Take a good look at Eq. 4.9. It doesn't make dimensional sense. If you ignore $|\Psi\rangle$ on both sides of the equation, the units on the left side are inverse time. If the quantum Hamiltonian is really to be identified with energy, then the units on the right side are energy. Energy is measured in units of joules, or $kg \cdot m^2/s^2$. Evidently, I've been cheating a little bit. The resolution of this dilemma involves \hbar , a universal constant of nature, which happens to have units of $kg \cdot m^2/s$. A constant with these units is exactly what we need to make Eq. 4.9 consistent. Let's rewrite it with Planck's constant inserted in a way that makes it dimensionally consistent:

$$\hbar \frac{\partial |\Psi\rangle}{\partial t} = -i\mathbf{H}|\Psi\rangle. \quad (4.10)$$

Why is it that \hbar is such a ridiculously small number? The answer has much more to do with biology than with physics.

The real question is not why \hbar is so small; it's why you are so big. The units that we use reflect our own size. The origin of the meter seems to be that it was used to measure rope or cloth: it's about the distance from a person's nose to his or her outstretched fingers. A second is about as long as a heartbeat. And a kilogram is a nice weight to carry around. We use these units because they are convenient, but fundamental physics doesn't care that much about us. The size of an atom is about 10^{-10} meters. Why so small? That's the wrong question. The right one is: Why are there so many atoms in an arm? The reason is simply that to make a functioning, intelligent, unit-using creature, you need to put together a lot of atoms. Similarly, the kilogram is many times larger than an atomic mass because people don't carry around single atoms; they get lost too easily. The same goes for time, and our long, plodding second. In the end, the reason that Planck's constant is so small is that we are so big and heavy and slow.

Physicists who are interested in the microscopic world are likely to use units that are more tailored to the phenomena that they study. If we used atomic length scales, time scales, and mass scales, then Planck's constant would not be such an unwieldy number; it would be much closer to 1. In fact, units for which Planck's constant equals 1 are a natural choice for quantum mechanics, and it's a common practice to use them. However, in this book, we will usually retain \hbar in our equations.

4.7 Expectation Values

Let's take a short break to discuss an important aspect of statistics, namely the idea of an average value or mean value. We mentioned this idea briefly in the previous lecture, but now it's time to take a closer look.

In quantum mechanics, average values are called expectation values. (In some ways, this is a poor choice of words; I'll tell you why later.) Suppose we have a probability function for the outcome of an experiment that measures an observable \mathbf{L} . The outcome must be one of \mathbf{L} 's eigenvalues, λ_i , and the probability function is $P(\lambda_i)$. In statistics, that average (or *mean*) value is denoted by a bar over the quantity being measured. The average of the observable \mathbf{L} would be $\bar{\mathbf{L}}$. In quantum mechanics, the standard notation is different, having grown out of Paul Dirac's clever bra-ket notation. We represent the average value of \mathbf{L} with the notation $\langle \mathbf{L} \rangle$. We'll soon see why the bra-ket notation is so natural, but first let's discuss the meaning of the term *average*.

From a mathematical point of view, an average is defined by the equation

$$\langle \mathbf{L} \rangle = \sum_i \lambda_i P(\lambda_i). \quad (4.11)$$

In other words, it is a weighted sum, weighted with the probability function P .

Alternatively, the average can be defined in an experimental way. Suppose a very large number of identical experiments is made, and the outcomes are recorded. Let's define

the probability function in a direct observational manner. We identify $P(\lambda_i)$ as the fraction of observations whose result was λ_i . The definition 4.11 is then identified with the experimental average of the observations. The basic hypothesis of any statistical theory is that if the number of trials is large enough, the mathematical and experimental notions of probability and average will agree. We will not question this hypothesis.

I'll now prove an elegant little theorem that explains the bra-ket notation for averages. Suppose that the normalized state of a quantum system is $|A\rangle$. Expand $|A\rangle$ in the orthonormal basis of eigenvectors of \mathbf{L} :

$$|A\rangle = \sum_i \alpha_i |\lambda_i\rangle. \quad (4.12)$$

Just for fun, with no particular agenda in mind, let's compute the quantity $\langle A|\mathbf{L}|A\rangle$. The meaning of this should be clear: First act on $|A\rangle$ with the linear operator \mathbf{L} .² Then, take the inner product of the result with the bra $\langle A|$. Let's do the first step by letting \mathbf{L} operate on both sides of Eq. 4.12:

$$\mathbf{L}|A\rangle = \sum_i \alpha_i \mathbf{L}|\lambda_i\rangle.$$

Remember that the vectors $|\lambda_i\rangle$ are eigenvectors of \mathbf{L} . Using the fact that $\mathbf{L}|\lambda_i\rangle = \lambda_i|\lambda_i\rangle$, we can write

²We would get the same result if we had let \mathbf{L} act on $\langle A|$ first.

$$\mathbf{L}|A\rangle = \sum_i \alpha_i \lambda_i |\lambda_i\rangle.$$

The last step is to take the inner product with $\langle A|$. We do that by expanding the bra $\langle A|$ in eigenvectors on the right-hand side, and then using the orthonormality of the eigenvectors. The result is

$$\langle A|\mathbf{L}|A\rangle = \sum_i (\alpha_i^* \alpha_i) \lambda_i. \quad (4.13)$$

Using the probability principle (Principle 4) to identify $(\alpha_i^* \alpha_i)$ with the probability $P(\lambda_i)$, we immediately see that the expression on the right side of Eq. 4.13 is the same as the expression on the right side of Eq. 4.11. That is to say,

$$\langle \mathbf{L} \rangle = \langle A|\mathbf{L}|A\rangle. \quad (4.14)$$

Thus, we have a quick rule to compute averages. Just sandwich the observable between the bra and ket representations of the state-vector.

In the previous lecture (Section 3.5), we promised to explain how the action of a Hermitian operator on a state-vector is related to the results of physical measurements. Armed with our knowledge of expectation values, we can now keep that promise. If we look back at Eq. 3.21, we see an example of an operator, σ_z , acting on state-vector $|r\rangle$ to produce a new state-vector. We can view this equation as half of the calculation for the expectation value of the measurement σ_z —the right-hand part of the sandwich, if you will.

The rest of that calculation involves taking the inner product of this state-vector with the dual vector $\langle r|$. So when σ_z acts on $|r\rangle$ in Eq. 3.21, it produces a state-vector from which we can calculate the probabilities of each σ_z measurement outcome.

4.8 Ignoring the Phase-Factor

In previous lectures, we said that we can ignore the overall phase-factor of a state-vector, and promised to explain why in a later section. Having worked out the rule for averages, we'll take a short detour to keep that promise.

What does it mean to “ignore the overall phase-factor”? It means we can multiply any state-vector by a constant factor $e^{i\theta}$, where θ is a real number, without changing the state-vector's physical meaning. To see this, let's multiply Eq. 4.12 by $e^{i\theta}$ and call the result $|B\rangle$:

$$|B\rangle = e^{i\theta}|A\rangle = e^{i\theta} \sum_j \alpha_j |\lambda_j\rangle. \quad (4.15)$$

Note that we changed the index in the summation from i to j to avoid confusion. It's easy to see that $|B\rangle$ has the same magnitude as $|A\rangle$, because $e^{i\theta}$ has a magnitude of one:

$$\langle B|B\rangle = \langle Ae^{-i\theta}|e^{i\theta}A\rangle = \langle A|A\rangle.$$

The same pattern of cancellation preserves other quantities as well. For example, $|A\rangle$'s probability amplitudes α_j become $e^{i\theta}\alpha_j$ for $|B\rangle$, so the probability *amplitudes* are different. However, it's the actual probability, not the amplitude,

that has physical meaning. If a system is in state $|B\rangle$, and we perform a measurement, the result will be the eigenvalue of $|\lambda_j\rangle$ with probability

$$\alpha_j^* e^{-i\theta} e^{i\theta} \alpha_j = \alpha_j^* \alpha_j,$$

which is the same result we would get for state $|A\rangle$. Finally, let's use the same trick for the expectation value of a Hermitian operator \mathbf{L} . Applying Eq. 4.14 to state $|B\rangle$, we can write

$$\langle \mathbf{L} \rangle = \langle B | \mathbf{L} | B \rangle.$$

Using Eq. 4.15 for $|B\rangle$, we get

$$\langle \mathbf{L} \rangle = \langle A e^{-i\theta} | \mathbf{L} | e^{i\theta} A \rangle$$

or

$$\langle \mathbf{L} \rangle = \langle A | \mathbf{L} | A \rangle.$$

In other words, \mathbf{L} has the same expectation value in state $|B\rangle$ as it does in state $|A\rangle$. Promise kept.

4.9 Connections to Classical Mechanics

The average, or expectation value, of an observable is the closest thing in quantum mechanics to a classical value. If

the probability distribution for an observable is a nice bell-shaped curve, and not too broad, then the expectation value really is the value that you expect to measure. If a system is so big and heavy that quantum mechanics is not too important, then the expectation value of an observable behaves almost exactly according to classical equations of motion. For this reason, it is interesting and important to find out how expectation values change with time.

First of all, *why* do they change with time? They change with time because the state of the system changes with time. Suppose the state at time t is represented by ket $|\Psi(t)\rangle$ and bra $\langle\Psi(t)|$. The expectation value of the observable \mathbf{L} at time t is

$$\langle\Psi(t)|\mathbf{L}|\Psi(t)\rangle.$$

Let's see how this changes by differentiating it with respect to t and using the Schrödinger equation for the time derivatives of $|\Psi(t)\rangle$ and $\langle\Psi(t)|$. Using the product rule for derivatives, we find that

$$\frac{d}{dt}\langle\Psi(t)|\mathbf{L}|\Psi(t)\rangle = \langle\dot{\Psi}(t)|\mathbf{L}|\Psi(t)\rangle + \langle\Psi(t)|\mathbf{L}|\dot{\Psi}(t)\rangle,$$

where, as usual, the dot means time derivative. \mathbf{L} itself has no explicit time dependency, so it just comes along for the ride. Now, plugging in the bra and ket versions of Schrödinger's equation (Eq. 4.10), we get

$$\frac{d}{dt}\langle\Psi(t)|\mathbf{L}|\Psi(t)\rangle = \frac{i}{\hbar}\langle\Psi(t)|\mathbf{H}\mathbf{L}|\Psi(t)\rangle - \frac{i}{\hbar}\langle\Psi(t)|\mathbf{L}\mathbf{H}|\Psi(t)\rangle$$

or, more concisely,

$$\frac{d}{dt}\langle\Psi(t)|\mathbf{L}|\Psi(t)\rangle = \frac{i}{\hbar}\langle\Psi(t)|[\mathbf{HL} - \mathbf{LH}]|\Psi(t)\rangle. \quad (4.16)$$

If you are used to ordinary algebra, Eq. 4.16 has a strange appearance. The right-hand side contains the combination $\mathbf{HL} - \mathbf{LH}$, a combination that would ordinarily be zero. But linear operators are not ordinary numbers: when they are multiplied (or applied sequentially), the order counts. In general, when \mathbf{H} acts on $\mathbf{L}|\Psi\rangle$, the result is not the same as when \mathbf{L} acts on $\mathbf{H}|\Psi\rangle$. In other words, except for special cases, $\mathbf{HL} \neq \mathbf{LH}$. Given two operators or matrices, the combination

$$\mathbf{LM} - \mathbf{ML}$$

is called the *commutator* of \mathbf{L} with \mathbf{M} , and it is denoted by a special symbol:

$$\mathbf{LM} - \mathbf{ML} = [\mathbf{L}, \mathbf{M}].$$

It's worth noticing that $[\mathbf{L}, \mathbf{M}] = -[\mathbf{M}, \mathbf{L}]$ for any pair of operators. Armed with the notation for commutators, we can now write Eq. 4.16 in a simple form:

$$\frac{d}{dt}\langle\mathbf{L}\rangle = \frac{i}{\hbar}\langle[\mathbf{H}, \mathbf{L}]\rangle \quad (4.17)$$

or, equivalently,

$$\frac{d}{dt}\langle\mathbf{L}\rangle = -\frac{i}{\hbar}\langle[\mathbf{L}, \mathbf{H}]\rangle. \quad (4.18)$$

This is a very interesting and important equation. It relates the time derivative of the expectation value of an observable \mathbf{L} to the expectation value of another observable, namely $-\frac{i}{\hbar}[\mathbf{L}, \mathbf{H}]$.

Exercise 4.2: Prove that if \mathbf{M} and \mathbf{L} are both Hermitian, $i[\mathbf{M}, \mathbf{L}]$ is also Hermitian. Note that the i is important. The commutator is, by itself, not Hermitian.

If we assume that the probabilities are nice, narrow, bell-shaped curves, then Eq. 4.18 tells us how the peaks of the curves move with time. Equations like this are the closest thing in quantum mechanics to the equations of classical physics. Sometimes we even omit the angle brackets in such equations and write them in a shorthand form:

$$\frac{d\mathbf{L}}{dt} = -\frac{i}{\hbar}[\mathbf{L}, \mathbf{H}]. \quad (4.19)$$

But keep in mind that a quantum equation of this type should be in the middle of a sandwich, with a bra $\langle\Psi|$ on one side, and a ket $|\Psi\rangle$ on the other. Alternatively, we can think of it as an equation that tells us how the centers of probability distributions move around.

Does Eq. 4.19 have a familiar look to it? If not, go back to Lectures 9 and 10 in *Volume I*, where we learned about the Poisson bracket formulation of classical mechanics. On

page 172, the following equation can be found:³

$$\dot{F} = \{F, H\}. \quad (4.20)$$

In this equation, $\{F, H\}$ is not a commutator; it is a Poisson bracket. But still, Eq. 4.20 is suspiciously similar to Eq. 4.19. In fact, there is a close parallel between commutators and Poisson brackets, and their algebraic properties are quite similar. For example, if F and G represent operators, both commutators and Poisson brackets change their sign when F and G are interchanged. Dirac discovered this, and realized that it represents an important structural connection between the mathematics of classical mechanics and that of quantum mechanics. The formal identification between commutators and Poisson brackets is

$$[\mathbf{F}, \mathbf{G}] \quad \iff \quad i\hbar\{F, G\}. \quad (4.21)$$

To facilitate comparison with Eq. 4.19, we can substitute the symbols \mathbf{L} and \mathbf{H} that we've been using in this section.

$$[\mathbf{L}, \mathbf{H}] \quad \iff \quad i\hbar\{L, H\}. \quad (4.22)$$

Let's try and make this identification as clear as possible. If we start with Eq. 4.19,

$$\frac{d\mathbf{L}}{dt} = -\frac{i}{\hbar}[\mathbf{L}, \mathbf{H}],$$

³ *Volume I*, Lecture 9, Eq. 10. Another one of those elegant French inventions.

and then use the identification of Eq. 4.22 to write the classical analog, the result is

$$\frac{d\mathbf{L}}{dt} = -\frac{i}{\hbar}(i\hbar\{\mathbf{L}, \mathbf{H}\})$$

or

$$\frac{d\mathbf{L}}{dt} = \{\mathbf{L}, \mathbf{H}\},$$

which matches the pattern of Eq. 4.20 exactly.

Exercise 4.3: Go back to the definition of Poisson brackets in *Volume I* and check that the identification in Eq. 4.21 is dimensionally consistent. Show that without the factor \hbar , it would not be.

Equation 4.21 solves a riddle. In classical physics, there is no difference between FG and GF . In other words: classically, commutators between ordinary observables are zero. From Eq. 4.21, we see that commutators in quantum mechanics are not zero, but that they are very small. The classical limit (the limit at which classical mechanics is accurate) is also the limit at which \hbar is negligibly small. Therefore, it is also the limit at which commutators are very small in human units.

4.10 Conservation of Energy

How can we tell whether something is conserved in quantum mechanics? What do we even mean by saying that an observable—call it Q —is conserved? At the very minimum,

we mean that its expectation value $\langle Q \rangle$ does not change with time (unless of course the system is disturbed). An even stronger condition is that $\langle Q^2 \rangle$ (or the expectation value of *any* power of Q) does not change with time.

Looking at Eq. 4.19, we can see that the condition for $\langle Q \rangle$ not to change is

$$[\mathbf{Q}, \mathbf{H}] = 0.$$

In other words, if a quantity commutes with the Hamiltonian, its expectation value is conserved. We can make this statement stronger. Using the properties of commutators, it's easy to see that if $[\mathbf{H}, \mathbf{Q}] = 0$, then $[\mathbf{Q}^2, \mathbf{H}] = 0$, or even more generally, $[\mathbf{Q}^n, \mathbf{H}] = 0$, for any n . It turns out that we can make a stronger claim: if \mathbf{Q} commutes with the Hamiltonian, the expectation values of *all* functions of \mathbf{Q} are conserved. That's what conservation means in quantum mechanics.

The most obvious conserved quantity is the Hamiltonian itself. Since any operator commutes with itself, one can write

$$[\mathbf{H}, \mathbf{H}] = 0,$$

which is exactly the condition that \mathbf{H} is conserved. As in classical mechanics, the Hamiltonian is another word for the energy of a system—it's a definition of energy. We see that under very general conditions, energy is conserved in quantum mechanics.

4.11 Spin in a Magnetic Field

Let's try out the Hamiltonian equations of motion for a single spin. We will first need to specify a Hamiltonian. Where do we get it from? In general, the answer is the same as in classical physics: derive it from experiment, borrow it from some theory that we like, or just pick one and see what it does. But in the case of a single spin, we don't have many options. Let's start with the unit operator I . Since I commutes with all operators, if it were the Hamiltonian, nothing would change with time. Remember, the time-dependence of an observable is given by the commutator of the observable with the Hamiltonian.

The only other choice is a sum of the spin components. In fact, that's exactly what we would get from experimental observation of a real spin—say an electron's spin—in a magnetic field. A magnetic field \vec{B} is a 3-vector—ordinary vector in space—and is specified by three Cartesian components, B_x , B_y , and B_z . When a classical spin (a charged rotor) is put into a magnetic field, it has an energy that depends on its orientation. The energy is proportional to the dot product of the spin and the magnetic field. The quantum version of this is

$$H \sim \vec{\sigma} \cdot \vec{B} = \sigma_x B_x + \sigma_y B_y + \sigma_z B_z,$$

where the symbol \sim means “proportional to.” Remember that σ_x , σ_y , and σ_z represent the components of the spin *operator* in the above quantum version.

Let's take a simple example in which the magnetic field lies along the z axis. In that case, the Hamiltonian is proportional to σ_z . For convenience, we'll absorb all the numerical constants, including the magnitude of the field (but not \hbar), into a single constant ω and write

$$\mathbf{H} = \frac{\hbar\omega}{2}\sigma_z. \quad (4.23)$$

The reason for the 2 in the denominator will become clear soon.

Our goal is to find out how the expectation value of the spin varies with time—in other words, to determine $\langle\sigma_x(t)\rangle$, $\langle\sigma_y(t)\rangle$, and $\langle\sigma_z(t)\rangle$. To do this, we just go back to Eq. 4.19, and plug in these components of \mathbf{L} . We get

$$\begin{aligned} \langle\dot{\sigma}_x\rangle &= -\frac{i}{\hbar}\langle[\sigma_x, \mathbf{H}]\rangle \\ \langle\dot{\sigma}_y\rangle &= -\frac{i}{\hbar}\langle[\sigma_y, \mathbf{H}]\rangle \\ \langle\dot{\sigma}_z\rangle &= -\frac{i}{\hbar}\langle[\sigma_z, \mathbf{H}]\rangle. \end{aligned} \quad (4.24)$$

Plugging in $\mathbf{H} = \frac{\hbar\omega}{2}\sigma_z$ from Eq. 4.23, we get

$$\begin{aligned} \langle\dot{\sigma}_x\rangle &= \frac{-i\omega}{2}\langle[\sigma_x, \sigma_z]\rangle \\ \langle\dot{\sigma}_y\rangle &= \frac{-i\omega}{2}\langle[\sigma_y, \sigma_z]\rangle \end{aligned}$$

$$\langle \dot{\sigma}_z \rangle = \frac{-i\omega}{2} \langle [\sigma_z, \sigma_z] \rangle. \quad (4.25)$$

The things we are computing on the left side of the equations are supposed to be real quantities. The factor i in these equations seems like trouble. Fortunately, the commutation relations between σ_x , σ_y , and σ_z will save the day. By plugging in the Pauli matrices from Eq. 3.20, it's easy to verify that

$$\begin{aligned} [\sigma_x, \sigma_y] &= 2i\sigma_z \\ [\sigma_y, \sigma_z] &= 2i\sigma_x \\ [\sigma_z, \sigma_x] &= 2i\sigma_y. \end{aligned} \quad (4.26)$$

Each of these equations also has an i , which will cancel the i in Eqs. 4.25. Notice that the factors of 2 also cancel, resulting in some very simple equations:

$$\begin{aligned} \langle \dot{\sigma}_x \rangle &= -\omega \langle \sigma_y \rangle \\ \langle \dot{\sigma}_y \rangle &= \omega \langle \sigma_x \rangle \\ \langle \dot{\sigma}_z \rangle &= 0. \end{aligned} \quad (4.27)$$

Does this look familiar? If not, go back to *Volume I*, Lecture 10. There, we studied the classical rotor in a magnetic field.

The equations were exactly the same, except that instead of expectation values, we were studying the actual motion of a deterministic system. Both there and here, the solution is that the 3-vector-operator $\vec{\sigma}$ (or the 3-vector \vec{L} in *Volume I*) precesses like a gyroscope around the direction of the magnetic field. The precession is uniform, with angular velocity ω .

This similarity to classical mechanics is very pleasing, but it's important to take note of the difference. Exactly *what* is precessing? In classical mechanics, it's just the x and y components of angular momentum. In quantum mechanics, it's an expectation value. The expectation value for a σ_z measurement does not change with time, but the other two expectation values do. Regardless, the result of each individual measurement of each spin component is still either $+1$ or -1 .

Exercise 4.4: Verify the commutation relations of Eqs. 4.26.

4.12 Solving the Schrödinger Equation

The iconic Schrödinger equation that appears on T-shirts has this form:

$$i\hbar \frac{\partial \Psi(x)}{\partial t} = -\frac{\hbar^2}{2m} \frac{\partial^2 \Psi(x)}{\partial x^2} + U(x)\Psi(x).$$

At this point, let's not worry about the meaning of the sym-

bols except to note that it is an equation that tells you how something changes with time. (The “something” is a representation of the state-vector of a particle.)

The iconic Schrödinger equation is a special case of a more general equation that we’ve already met in Eq. 4.9. It is part definition and part principle of quantum mechanics. As a principle, it says that the state-vector changes continuously with time, in a unitary way. As a definition, it defines the Hamiltonian, and therefore the observable called energy. Eq. 4.10,

$$\hbar \frac{\partial |\Psi\rangle}{\partial t} = -i\mathbf{H}|\Psi\rangle,$$

is sometimes called the *time-dependent* Schrödinger equation. Because the Hamiltonian operator \mathbf{H} represents energy, the observable values of energy are just the eigenvalues of \mathbf{H} . Let’s call these eigenvalues E_j and the corresponding eigenvectors $|E_j\rangle$. By definition, the relation between \mathbf{H} , E_j , and $|E_j\rangle$ is the eigenvalue equation

$$\mathbf{H}|E_j\rangle = E_j|E_j\rangle. \quad (4.28)$$

This is the *time-independent Schrödinger equation*, and it’s used in two different ways.

If we work in a particular matrix basis, then the equation determines the eigenvectors of \mathbf{H} . One puts in a particular value of the energy E_j and looks for the ket-vector $|E_j\rangle$ that solves the equation.

It is also an equation that determines the eigenvalues E_j . If you put in an arbitrary value of E_j , in general there will

not be a solution for the eigenvector. Let's take a very simple example: Suppose the Hamiltonian is the matrix $\frac{\hbar\omega}{2}\sigma_z$. Since σ_z has only two eigenvalues, namely ± 1 , the Hamiltonian also has only two eigenvalues, $\pm\frac{\hbar\omega}{2}$. If you put any other value on the right-hand side of Eq. 4.28, there will not be a solution. Because the operator \mathbf{H} represents energy, we often call E_j the *energy eigenvalues* and $|E_j\rangle$ the *energy eigenvectors* of the system.

Exercise 4.5: Take any unit 3-vector \vec{n} and form the operator

$$\mathbf{H} = \frac{\hbar\omega}{2}\sigma \cdot \vec{n}.$$

Find the energy eigenvalues and eigenvectors by solving the time-independent Schrödinger equation. Recall that Eq. 3.23 gives $\sigma \cdot \vec{n}$ in component form.

Let's suppose we have found all the energy eigenvalues E_j and the corresponding eigenvectors $|E_j\rangle$. We can now use that information to solve the time-dependent Schrödinger equation. The trick is to use the fact that the eigenvectors form an orthonormal basis and then expand the state-vector in that basis. Let the state-vector be called $|\Psi\rangle$ and write

$$|\Psi\rangle = \sum_j \alpha_j |E_j\rangle.$$

Since the state-vector $|\Psi\rangle$ changes with time and the basis vectors $|E_j\rangle$ do not, it follows that the coefficients α_j must

also depend on time:

$$|\Psi(t)\rangle = \sum_j \alpha_j(t) |E_j\rangle. \quad (4.29)$$

Now feed Eq. 4.29 into the time-dependent equation. The result is

$$\sum_j \dot{\alpha}_j(t) |E_j\rangle = -\frac{i}{\hbar} \mathbf{H} \sum_j \alpha_j(t) |E_j\rangle.$$

Next, we use the fact that $\mathbf{H}|E_j\rangle = E_j|E_j\rangle$ to get

$$\sum_j \dot{\alpha}_j(t) |E_j\rangle = -\frac{i}{\hbar} \sum_j E_j \alpha_j(t) |E_j\rangle$$

or, regrouping,

$$\sum_j \left\{ \dot{\alpha}_j(t) + \frac{i}{\hbar} E_j \alpha_j(t) \right\} |E_j\rangle = 0.$$

The final step should be easy to see. If a sum of basis vectors equals zero, every coefficient must be zero. Hence, for each eigenvalue E_j , $\alpha_j(t)$ must satisfy the simple differential equation

$$\frac{d\alpha_j(t)}{dt} = -\frac{i}{\hbar} E_j \alpha_j(t).$$

This, of course, is the familiar differential equation for an exponential function of time, in this case with an imaginary exponent. The solution is

$$\alpha_j(t) = \alpha_j(0) e^{-\frac{i}{\hbar} E_j t}. \quad (4.30)$$

This equation tells us how the α_j change with time. It is quite general and not restricted to spins, provided that the Hamiltonian does not depend explicitly on time. This is our first example of the deep connection between energy and frequency, which recurs over and over throughout quantum mechanics and quantum field theory. We will return to it often.

In Eq. 4.30, the factors $\alpha_j(0)$ are the values of the coefficients at time zero. If we know the state-vector $|\Psi\rangle$ at time zero, then the coefficients are given by the projections of $|\Psi\rangle$ on the basis eigenvectors. We can write this as

$$\alpha_j(0) = \langle E_j | \Psi(0) \rangle. \quad (4.31)$$

Now let's put the whole thing together and write the full solution of the time-dependent Schrödinger equation:

$$|\Psi(t)\rangle = \sum_j \alpha_j(0) e^{-\frac{i}{\hbar} E_j t} |E_j\rangle.$$

When we use Eq. 4.31 to replace $\alpha_j(0)$, this equation becomes

$$|\Psi(t)\rangle = \sum_j \langle E_j | \Psi(0) \rangle e^{-\frac{i}{\hbar} E_j t} |E_j\rangle. \quad (4.32)$$

Eq. 4.32 can be written in the more elegant form,

$$|\Psi(t)\rangle = \sum_j |E_j\rangle \langle E_j | \Psi(0) \rangle e^{-\frac{i}{\hbar} E_j t}, \quad (4.33)$$

which emphasizes that we're summing over the basis vectors. You may wonder how we just happen to “know” $|\Psi(0)\rangle$. The answer depends on the circumstances, but usually, we assume we can use some apparatus to prepare the system in a known state.

Before we discuss the bigger meaning of these equations, I want to restate them as a recipe. I'll assume you already know enough about the system and its space of states to get started.

4.13 Recipe for a Schrödinger Ket

1. Derive, look up, guess, borrow, or steal the Hamiltonian operator \mathbf{H} .
2. Prepare an initial state $|\Psi(0)\rangle$.
3. Find the eigenvalues and eigenvectors of \mathbf{H} by solving the time-independent Schrödinger equation,

$$\mathbf{H}|E_j\rangle = E_j|E_j\rangle.$$

4. Use the initial state-vector $|\Psi(0)\rangle$, along with the eigenvectors $|E_j\rangle$ from step 3, to calculate the initial coefficients $\alpha_j(0)$:

$$\alpha_j(0) = \langle E_j|\Psi(0)\rangle.$$

5. Rewrite $|\Psi(0)\rangle$ in terms of the eigenvectors $|E_j\rangle$ and the initial coefficients $\alpha_j(0)$:

$$|\Psi(0)\rangle = \sum_j \alpha_j(0)|E_j\rangle.$$

What we've done so far is to expand the initial state-vector $|\Psi(0)\rangle$ in terms of the eigenvectors $|E_j\rangle$ of \mathbf{H} . Why is that basis better than any other? Because \mathbf{H} tells us how things evolve with time. We will use that knowledge now.

6. In the above equation, replace each $\alpha_j(0)$ with $\alpha_j(t)$ to capture its time-dependence. As a result, $|\Psi(0)\rangle$ becomes $|\Psi(t)\rangle$:

$$|\Psi(t)\rangle = \sum_j \alpha_j(t) |E_j\rangle.$$

7. Using Eq. 4.30, replace each $\alpha_j(t)$ with $\alpha_j(0)e^{-\frac{i}{\hbar}E_j t}$:

$$|\Psi(t)\rangle = \sum_j \alpha_j(0) e^{-\frac{i}{\hbar}E_j t} |E_j\rangle. \quad (4.34)$$

8. Season according to taste.

We can now predict the probabilities for each possible outcome of an experiment as a function of time, and we are not restricted to energy measurements. Suppose \mathbf{L} has eigenvalues λ_j and eigenvectors $|\lambda_j\rangle$. The probability for outcome λ is

$$P_\lambda(t) = |\langle \lambda | \Psi(t) \rangle|^2.$$

Exercise 4.6: Carry out the Schrödinger Ket recipe for a single spin. The Hamiltonian is $\mathbf{H} = \frac{\omega\hbar}{2}\sigma_z$ and the final observable is σ_x . The initial state is given as $|u\rangle$ (the state in which $\sigma_z = +1$).

After time t , an experiment is done to measure σ_y . What are the possible outcomes and what are the probabilities for those outcomes?

Congratulations! You have now solved a real quantum mechanics problem for an experiment that can actually be carried out in the laboratory. Feel free to pat yourself on the back.

4.14 Collapse

We've seen how the state-vector evolves between the time that a system is prepared in a given state and the time that it is brought into contact with an apparatus and measured. If the state-vector were main focus of observational physics, we would say that quantum mechanics is deterministic. But experimental physics is not about measuring the state-vector. It is about measuring observables. Even if we know the state-vector exactly, we don't know the result of any given measurement. Nevertheless, it is fair to say that between observations, the state of a system evolves in a perfectly definite way, according to the time-dependent Schrödinger equation.

But something different happens when an observation is made. An experiment to measure \mathbf{L} will have an unpredictable outcome, but after the measurement is made, the system is left in an eigenstate of \mathbf{L} . Which eigenstate? The

one corresponding to the outcome of the measurement. But this outcome is unpredictable. So it follows that during an experiment the state of a system jumps unpredictably to an eigenstate of the observable that was measured. This phenomenon is called *the collapse of the wave function*.⁴

To put it another way, suppose the state-vector is

$$\sum_j \alpha_j |\lambda_j\rangle$$

just before the measurement of \mathbf{L} . Randomly, with probability $|\alpha_j|^2$, the apparatus measures a value λ_j and leaves the system in a single eigenstate of \mathbf{L} , namely $|\lambda_j\rangle$. The entire superposition of states collapses to a single term.

This strange fact—that the system evolves one way between measurements and another way during a measurement—has been a source of contention and confusion for decades. It raises a question: Shouldn't the act of measurement itself be described by the laws of quantum mechanics?

The answer is yes. The laws of quantum mechanics are not suspended during measurement. However, to examine the measurement process itself as a quantum mechanical evolution, we must consider the entire experimental setup, including the apparatus, as part of a single quantum system. We'll discuss that topic—how systems are combined into composite systems—in Lecture 6. But first, a few words about uncertainty.

⁴We have not yet explained what a wave function is, but we'll do so shortly, in Section 5.1.2.

Lecture 5

Uncertainty and Time Dependence

Lenny: *Good evening, General. Nice to see you again.*

The General: *Lenny? Is that you? It's been forever. Well, a long time anyway. Who's your friend?*

Lenny: *His name is Art. Art, shake hands with General Uncertainty.*

5.1 Mathematical Interlude: Complete Sets of Commuting Variables

5.1.1 States That Depend On More Than One Measurable

The physics of a single spin is extremely simple, and that's what makes it so attractive as an illustrative example. But

that also means there's a lot it can't illustrate. One property of a single spin is that its state can be fully specified by the eigenvalue of a single operator, say σ_z . If the value of σ_z is known, then no other observable—such as σ_x —can also be specified. As we have seen, measuring either of these quantities destroys any information we may have had about the other one.

But in more complicated systems, we may have multiple observables that are compatible; that is, their values can be known simultaneously. Here are two examples:

- A particle moving in three-dimensional space. A basis of states for this system is specified by the position of the particle, but this takes three position coordinates. Thus, we have states that are specified by three numbers, $|x, y, z\rangle$. We will see later that all three spatial coordinates of a particle can be simultaneously specified.
- A system composed of two physically independent spins; in other words, a system of two qubits. Later, we will see how to combine systems to form bigger systems. But for now we can just say that the two-spin system can be described by two observables. Namely, we have a state in which both spins are up, another in which both are down, another in which the first is up while the second is down, and another in which these spins are reversed. To put it more briefly, we can characterize the two-spin system by two observables: the z component of the first spin and the z component of the second spin. Quantum mechanics does not forbid

simultaneous knowledge of these two observables. In fact, one can choose any component of one spin and any component of the other spin. Quantum mechanics allows simultaneous knowledge of both.

In these situations, we need multiple measurements to fully characterize the state of the system. For example, in our two-spin system, we measure each spin separately and associate these measurements with two different operators. We'll call these operators \mathbf{L} and \mathbf{M} .

A measurement leaves the system in an eigenstate (consisting of a single eigenvector), corresponding to the value (an eigenvalue) that was measured. If we measure both spins in a two-spin system, the system winds up in a state that is simultaneously an eigenvector of \mathbf{L} and an eigenvector of \mathbf{M} . We call this a *simultaneous eigenvector* of the operators \mathbf{L} and \mathbf{M} .

The two-spin example gives us something concrete to think about, but keep in mind that our results will be far more general—they will apply to any system that is characterized by two different operators. And as you might guess, there is nothing magic about the number two. The ideas presented here generalize to larger systems that require many operators to characterize them.

To work with two different compatible operators, we'll need two sets of labels for their basis vectors. We'll use the labels λ_i and μ_a . The symbols λ_i and μ_a are the eigenvalues of \mathbf{L} and \mathbf{M} . The subscripts i and a run over all the possible outcomes of measurements of \mathbf{L} and \mathbf{M} . We assume that

there is a basis of state-vectors $|\lambda_i, \mu_a\rangle$ that are simultaneous eigenvectors of both observables. In other words,

$$\mathbf{L}|\lambda_i, \mu_a\rangle = \lambda_i|\lambda_i, \mu_a\rangle$$

$$\mathbf{M}|\lambda_i, \mu_a\rangle = \mu_a|\lambda_i, \mu_a\rangle.$$

To make these equations a little less precise but a little easier to read, I will sometimes leave out the subscripts:

$$\mathbf{L}|\lambda, \mu\rangle = \lambda|\lambda, \mu\rangle$$

$$\mathbf{M}|\lambda, \mu\rangle = \mu|\lambda, \mu\rangle.$$

In order to have a basis of simultaneous eigenvectors, the operators \mathbf{L} and \mathbf{M} must commute. This is easy to see. We begin by acting on any of the basis vectors with the product \mathbf{LM} , and then use the fact that the basis vector is an eigenvector of both:

$$\mathbf{LM}|\lambda, \mu\rangle = \mathbf{L}\mu|\lambda, \mu\rangle,$$

or

$$\mathbf{LM}|\lambda, \mu\rangle = \lambda\mu|\lambda, \mu\rangle.$$

The eigenvalues λ , μ are of course just numbers and it doesn't matter which one appears first when we multiply them. Thus, if we reverse the order of these operators, and let the operator \mathbf{ML} act on the same basis vector, we get the same result:

$$\mathbf{LM}|\lambda, \mu\rangle = \mathbf{ML}|\lambda, \mu\rangle,$$

or, more succinctly,

$$[\mathbf{L}, \mathbf{M}] |\lambda, \mu\rangle = 0, \quad (5.1)$$

where the right-hand side represents the zero vector. This result would not be very helpful if it were only true for a particular basis vector. But the reasoning that leads us to Eq. 5.1 is valid for *any* of the basis vectors. That's enough to ensure that the operator $[\mathbf{L}, \mathbf{M}] = 0$. If an operator *annihilates* every member of a basis, it must also annihilate every vector in the vector space.¹ An operator that annihilates every vector is exactly what we mean by the *zero operator*. Thus, we prove that if there is a complete basis of simultaneous eigenvectors of two observables, the two observables must commute. It turns out that the converse of this theorem is also true: if two observables commute, then there is a complete basis of simultaneous eigenvectors of the two observables. To put it simply, the condition for two observables to be simultaneously measurable is that they commute.

As we mentioned earlier, this theorem is more general. One may need to specify a larger number of observables to completely label a basis. Regardless of the number of observables that are needed, they must all commute among themselves. We call this collection a *complete set of commuting observables*.

¹Do you see why?

5.1.2 Wave Functions

Now we'll introduce a concept called the *wave function*. For now, ignore the name; in general, the quantum wave function may have nothing to do with waves. Later, when we study the quantum mechanics of particles (Lectures 8–10), we'll find out about the connection between wave functions and waves.

Suppose we have a basis of states for some quantum system. Let the orthonormal basis vectors be called $|a, b, c, \dots\rangle$, where a, b, c, \dots are the eigenvalues of some complete set of commuting observables $\mathbf{A}, \mathbf{B}, \mathbf{C}, \dots$. Now, consider an arbitrary state vector $|\Psi\rangle$. Since the vectors $|a, b, c, \dots\rangle$ are an orthonormal basis, $|\Psi\rangle$ can be expanded in terms of them:

$$|\Psi\rangle = \sum_{a,b,c,\dots} \psi(a, b, c, \dots) |a, b, c, \dots\rangle.$$

The quantities $\psi(a, b, c, \dots)$ are the coefficients that enter the expansion. Each of them is also equal to the inner product of $|\Psi\rangle$ with one of the basis vectors:

$$\psi(a, b, c, \dots) = \langle a, b, c, \dots | \Psi \rangle. \quad (5.2)$$

The set of coefficients $\psi(a, b, c, \dots)$ is called the wave function of the system in the basis defined by the observables $\mathbf{A}, \mathbf{B}, \mathbf{C}, \dots$. The mathematical definition of a wave function is given by Eq. 5.2, which seems formal and abstract, but the physical meaning of the wave function is profoundly important. According to the basic probability principle of quantum mechanics, the squared magnitude of the wave function

is the probability for the commuting observables to have values a, b, c, \dots :

$$P(a, b, c, \dots) = \psi^*(a, b, c, \dots) \psi(a, b, c, \dots).$$

The form of the wave function depends on which observables we choose to focus on. That's because calculations for two different observables rely on different sets of basis vectors. For example, in the case of a single spin, the inner products

$$\psi(u) = \langle u | \Psi \rangle$$

and

$$\psi(d) = \langle d | \Psi \rangle$$

define the wave function in the σ_z basis, while

$$\psi(r) = \langle r | \Psi \rangle$$

and

$$\psi(l) = \langle l | \Psi \rangle$$

define the wave function in the σ_x basis.

An important feature of the wave function follows from the fact that the total probability sums to one:

$$\sum_{a,b,c,\dots} \psi^*(a, b, c, \dots) \psi(a, b, c, \dots) = 1.$$

5.1.3 A Note About Terminology

The term *wave function*, as used in this book, refers to the collection of coefficients (also called components) that multiply the basis vectors in an eigenfunction expansion. For example, if we expand a state-vector $|\Psi\rangle$ as follows,

$$|\Psi\rangle = \sum_j \alpha_j |\psi_j\rangle,$$

where the $|\psi_j\rangle$ are the orthonormal eigenvectors of a Hermitian operator, the collection of coefficients α_j —the things we called $\psi(a, b, c, \dots)$ just above—is what we mean by the wave function. In situations where the state-vector is expressed as an integral rather than a sum, the wave function is continuous rather than discrete.

So far, we have been careful to distinguish the wave function from the state-vectors $|\psi_j\rangle$, and this is a common convention. However, some authors refer to wave functions as though *they* are the state-vectors. This ambiguous use of terminology can be confusing. It becomes less confusing when you realize that a wave function really can represent a state-vector. It is reasonable to think of the α_j coefficients as the coordinates of the state-vector in a specific basis of eigenvectors. This is similar to saying that a set of Cartesian coordinates represents a particular point in 3-space relative to a specific coordinate frame. To avoid confusion, just try to be aware of which convention is being followed. In this book, we will generally use uppercase symbols, such as Ψ , to represent state-vectors, and lowercase symbols, such as ψ , to represent wave functions.

5.2 Measurement

Let's return to the concept of measurement. Suppose we measure two observables \mathbf{L} and \mathbf{M} in a single experiment, and the system is left in a simultaneous eigenvector of these two observables. As we learned in Section 5.1.1, this means that \mathbf{L} and \mathbf{M} must commute.

But what if they don't commute? Then, in general, it is not possible to have unambiguous knowledge of both. Later on, we will make this more quantitative in the form of the uncertainty principle, Heisenberg's being a special case.

Let's go back to our touchstone, the problem of a single spin. Any observable of a spin is represented by a 2×2 Hermitian matrix, and any such matrix has the form

$$\begin{pmatrix} r & w \\ w^* & r' \end{pmatrix},$$

with the diagonal elements being real and the other two being complex conjugates. The implication is that it takes exactly four real parameters to specify this observable. In fact, there is a neat way to write any spin observable in terms of the Pauli matrices, σ_x , σ_y , and σ_z , and one more matrix: the unit matrix I . As you recall,

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

$$\sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}$$

$$\sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

$$I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

Any 2×2 Hermitian matrix \mathbf{L} can be written as a sum of four terms,

$$\mathbf{L} = a\sigma_x + b\sigma_y + c\sigma_z + dI,$$

where a , b , c , and d are real numbers.

Exercise 5.1: Verify this claim.

The unit operator I is officially an observable because it is Hermitian, but it's a very boring one. There is only one possible value this trivial observable can have, namely 1, and every state-vector is an eigenvector. If we ignore I , then the most general observable is a superposition of the three spin components σ_x , σ_y , and σ_z . Can any pair of spin components be simultaneously measured? Only if they commute. But it is easy to calculate the commutators for these spin components. Just use the matrix representation to multiply them in both orders, and then subtract.

The commutation relations we listed in Eqs. 4.26,

$$[\sigma_x, \sigma_y] = 2i\sigma_z$$

$$\begin{aligned}[\sigma_y, \sigma_z] &= 2i\sigma_x \\ [\sigma_z, \sigma_x] &= 2i\sigma_y,\end{aligned}$$

tell us straightaway that no two spin components can be simultaneously measured, because the right-hand sides are not zero. In fact, no two components of the spin along any axes can be simultaneously measured.

5.3 The Uncertainty Principle

Uncertainty is one of the hallmarks of quantum mechanics, but it is not always the case that the result of an experiment is uncertain. If a system is in an eigenstate of an observable, then there is no uncertainty about the result of measuring that observable. But whatever the state, there is always uncertainty about some observable. If the state happens to be an eigenvector of one Hermitian operator—call it \mathbf{A} —then it will not be an eigenvector of other operators that don't commute with \mathbf{A} . Thus, as a rule, if \mathbf{A} and \mathbf{B} do not commute, then there must be uncertainty in one or the other, if not both.

The iconic example of this mutual uncertainty is the Heisenberg Uncertainty Principle, which in its original form had to do with the position and momentum of a particle. But Heisenberg's ideas can be expanded into a much more general principle that applies to any two observables that happen not to commute. An example would be two components of a spin. We now have all the ingredients necessary to derive the general form of the uncertainty principle.

5.4 The Meaning of Uncertainty

We need to be very certain about what we mean by *uncertainty* if we want to quantify it. Let's suppose the eigenvalues of the observable \mathbf{A} are called a . Then, given a state $|\Psi\rangle$, there is a probability distribution $P(a)$ with the usual properties. The expectation value of \mathbf{A} is the ordinary average:

$$\langle\Psi|\mathbf{A}|\Psi\rangle = \sum_a aP(a).$$

Roughly speaking, this means that $P(a)$ is centered around the expectation value. What we will mean by “the uncertainty in \mathbf{A} ” is the so-called *standard deviation*. To compute the standard deviation, begin by subtracting from \mathbf{A} its expectation value. We define the operator $\bar{\mathbf{A}}$ to be:

$$\bar{\mathbf{A}} = \mathbf{A} - \langle\mathbf{A}\rangle.$$

By defining $\bar{\mathbf{A}}$ in this way, we have subtracted an expectation value from an operator, and it's not completely clear what that means. Let's take a closer look. The expectation value is itself a real number. Every real number is also an operator, namely an operator proportional to the identity or unit operator I . To make the meaning clear, we can write $\bar{\mathbf{A}}$ in a more complete form:

$$\bar{\mathbf{A}} = \mathbf{A} - \langle\mathbf{A}\rangle I.$$

The probability distribution for $\bar{\mathbf{A}}$ is exactly the same as the distribution for \mathbf{A} except that it is shifted so that the average

of $\bar{\mathbf{A}}$ is zero. The eigenvectors of $\bar{\mathbf{A}}$ are the same as those of \mathbf{A} and the eigenvalues are just shifted so that their average is zero as well. In other words, the eigenvalues of $\bar{\mathbf{A}}$ are

$$\bar{a} = a - \langle \mathbf{A} \rangle.$$

The square of the uncertainty (or standard deviation) of \mathbf{A} , which we call $(\Delta \mathbf{A})^2$, is defined by

$$(\Delta \mathbf{A})^2 = \sum_a \bar{a}^2 P(a) \quad (5.3)$$

or

$$(\Delta \mathbf{A})^2 = \sum_a (a - \langle \mathbf{A} \rangle)^2 P(a). \quad (5.4)$$

This may also be written as

$$(\Delta \mathbf{A})^2 = \langle \Psi | \bar{\mathbf{A}}^2 | \Psi \rangle.$$

If the expectation value of \mathbf{A} is zero, then the uncertainty $\Delta \mathbf{A}$ takes the simpler form

$$(\Delta \mathbf{A})^2 = \langle \Psi | \mathbf{A}^2 | \Psi \rangle.$$

In other words the square of the uncertainty is the average value of the operator \mathbf{A}^2 .

5.5 Cauchy-Schwarz Inequality

The uncertainty principle is an inequality that says the product of the uncertainties of \mathbf{A} and \mathbf{B} is larger than something that involves their commutator. The basic mathematical inequality is the familiar triangle inequality. It says that in any vector space, the magnitude of one side of a triangle is less than the sum of the magnitudes of the other two sides. For real vector spaces, we derive

$$|X||Y| \geq |X \cdot Y| \quad (5.5)$$

from the triangle inequality,

$$|X| + |Y| \geq |X + Y|.$$

5.6 The Triangle Inequality and the Cauchy-Schwarz Inequality

The triangle inequality is motivated, of course, by the properties of ordinary triangles, but it's actually far more general and applies to a large class of vector spaces. You can get the basic idea by looking at Fig. 5.1, where the sides of the triangle are taken to be ordinary geometric vectors in a plane. The triangle inequality is just the statement that the sum of any two sides is bigger than the third side, and the underlying idea is that the shortest path between two points is a straight line. The shortest path between point 1 and point 3 is side Z , and the sum of the other two sides is certainly bigger.

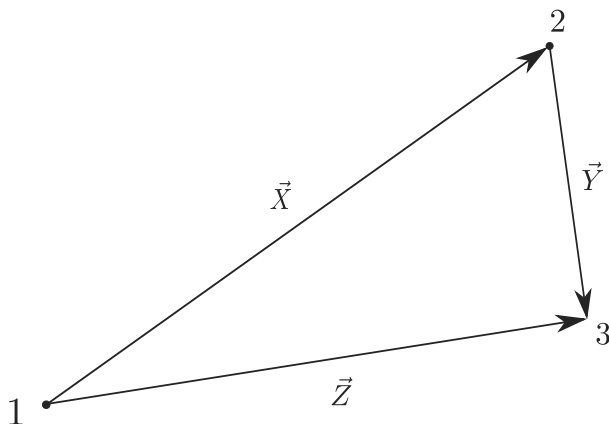


Figure 5.1: The Triangle Inequality. The sum of the lengths of vectors \vec{X} and \vec{Y} is greater than or equal to the length of vector \vec{Z} . (The shortest path between two points is a straight line.)

The triangle inequality can be expressed in more than one way. We'll start with the basic definition and then massage it into the form we need. We know that

$$|X| + |Y| \geq |Z|.$$

If we think of X and Y as vectors that can be added, we can write the above as

$$|\vec{X}| + |\vec{Y}| \geq |\vec{X} + \vec{Y}|.$$

If we square this equation, it becomes

$$|\vec{X}|^2 + |\vec{Y}|^2 + 2|\vec{X}||\vec{Y}| \geq |\vec{X} + \vec{Y}|^2.$$

But the right-hand side can be expanded as

$$|\vec{X} + \vec{Y}|^2 = |\vec{X}|^2 + |\vec{Y}|^2 + 2(\vec{X} \cdot \vec{Y}).$$

Why? Because $|\vec{X} + \vec{Y}|^2$ is just $(\vec{X} + \vec{Y}) \cdot (\vec{X} + \vec{Y})$. Collecting these results, we get

$$|\vec{X}|^2 + |\vec{Y}|^2 + 2|\vec{X}||\vec{Y}| \geq |\vec{X}|^2 + |\vec{Y}|^2 + 2(\vec{X} \cdot \vec{Y}).$$

Now, we just subtract $|\vec{X}|^2 + |\vec{Y}|^2$ from each side and then divide by 2, leaving us with

$$|\vec{X}||\vec{Y}| \geq \vec{X} \cdot \vec{Y}. \quad (5.6)$$

This is another form of the triangle inequality. It says that, given any two vectors \vec{X} and \vec{Y} , the product of their lengths is greater than or equal to their dot product. This is no surprise—the dot product is often defined as

$$\vec{X} \cdot \vec{Y} = |\vec{X}||\vec{Y}| \cos \theta,$$

where θ is the angle between the two vectors. But we know that the cosine of an angle always stays in the range -1 to $+1$, so the right-hand side must always be less than or equal to $|\vec{X}||\vec{Y}|$. This relationship is true for vectors in two dimensions, three dimensions, or an arbitrary number of dimensions. It's even true for vectors in complex vector spaces. It's generally true for vectors in *any* vector space, provided

the length of the vector is defined as the square root of the vector's inner product with itself. As we go forward, we plan to use Inequality 5.6 in the squared form, that is,

$$|\vec{X}|^2|\vec{Y}|^2 \geq (\vec{X} \cdot \vec{Y})^2$$

or

$$|\vec{X}|^2|\vec{Y}|^2 \geq |\vec{X} \cdot \vec{Y}|^2. \quad (5.7)$$

In this form, it's called the *Cauchy-Schwarz* inequality.

For complex vector spaces, the triangle inequality takes a slightly more complicated form. Let $|X\rangle$ and $|Y\rangle$ be any two vectors in a complex vector space. The magnitudes of the three vectors $|X\rangle$, $|Y\rangle$, and $|X\rangle + |Y\rangle$ are

$$\begin{aligned} |X| &= \sqrt{\langle X|X\rangle} \\ |Y| &= \sqrt{\langle Y|Y\rangle} \\ |X + Y| &= \sqrt{(\langle X| + \langle Y|)(|X\rangle + |Y\rangle)} \end{aligned} \quad (5.8)$$

We now follow the same steps as we did for the real case: First write

$$|X| + |Y| \geq |X + Y|.$$

Then square it and simplify:

$$2|X||Y| \geq |\langle X|Y\rangle + \langle Y|X\rangle|. \quad (5.9)$$

This is the form of the Cauchy-Schwarz inequality that will lead to the uncertainty principle. But what does it have to do with the two observables \mathbf{A} and \mathbf{B} ? We'll find out by cleverly defining $|X\rangle$ and $|Y\rangle$.

5.7 The General Uncertainty Principle

Let $|\Psi\rangle$ be any ket and let \mathbf{A} and \mathbf{B} be any two observables. We now define $|X\rangle$ and $|Y\rangle$ as follows:

$$\begin{aligned} |X\rangle &= \mathbf{A}|\Psi\rangle \\ |Y\rangle &= i\mathbf{B}|\Psi\rangle. \end{aligned} \quad (5.10)$$

Notice the i in the second definition. Now, substitute 5.10 into 5.9 to get

$$2\sqrt{\langle \mathbf{A}^2 \rangle \langle \mathbf{B}^2 \rangle} \geq |\langle \Psi | \mathbf{A} \mathbf{B} | \Psi \rangle - \langle \Psi | \mathbf{B} \mathbf{A} | \Psi \rangle|. \quad (5.11)$$

The minus sign is due to the factor of i in the second definition in 5.10. Using the definition of a commutator, we find that

$$2\sqrt{\langle \mathbf{A}^2 \rangle \langle \mathbf{B}^2 \rangle} \geq |\langle \Psi | [\mathbf{A}, \mathbf{B}] | \Psi \rangle|. \quad (5.12)$$

Let's suppose for the moment that \mathbf{A} and \mathbf{B} have expectation values of zero. In that case, $\langle \mathbf{A}^2 \rangle$ is just the square of the uncertainty in \mathbf{A} , that is, $(\Delta \mathbf{A})^2$, and $\langle \mathbf{B}^2 \rangle$ is just $(\Delta \mathbf{B})^2$. Thus we can rewrite Eq. 5.12 as

$$\Delta \mathbf{A} \Delta \mathbf{B} \geq \frac{1}{2} |\langle \Psi | [\mathbf{A}, \mathbf{B}] | \Psi \rangle|. \quad (5.13)$$

Reflect on this mathematical inequality for a moment. On the left side, we see the product of the uncertainties of the two observables \mathbf{A} and \mathbf{B} in the state Ψ . The inequality says that this product cannot be smaller than the right side, which involves the commutator of \mathbf{A} and \mathbf{B} . Specifically, it says that the product of the uncertainties cannot be smaller than *half the magnitude of the expectation value of the commutator*.

The general uncertainty principle is a quantitative expression of something we already suspected: if the commutator of \mathbf{A} and \mathbf{B} is not zero, then both observables cannot simultaneously be certain.

But what if the expectation value of \mathbf{A} or \mathbf{B} is not zero? In that case, the trick is to redefine two new operators in which the expectation values have been subtracted off:

$$\bar{\mathbf{A}} = \mathbf{A} - \langle \mathbf{A} \rangle$$

$$\bar{\mathbf{B}} = \mathbf{B} - \langle \mathbf{B} \rangle.$$

Then repeat the whole process, replacing \mathbf{A} and \mathbf{B} with $\bar{\mathbf{A}}$ and $\bar{\mathbf{B}}$. The following exercise serves as a guide.

Exercise 5.2:

- 1) Show that $\Delta \mathbf{A}^2 = \langle \bar{\mathbf{A}}^2 \rangle$ and $\Delta \mathbf{B}^2 = \langle \bar{\mathbf{B}}^2 \rangle$.
- 2) Show that $[\bar{\mathbf{A}}, \bar{\mathbf{B}}] = [\mathbf{A}, \mathbf{B}]$.
- 3) Using these relations, show that

$$\Delta \mathbf{A} \Delta \mathbf{B} \geq \frac{1}{2} |\langle \Psi | [\mathbf{A}, \mathbf{B}] | \Psi \rangle|.$$

Later, in Lecture 8, we will use this very general version of the uncertainty principle to prove the original form of Heisenberg's Uncertainty Principle: *The product of the uncertainties of the position and momentum of a particle cannot be less than half of Planck's constant.*

Lecture 6

Combining Systems: Entanglement

Art: *This is a pretty friendly place after all. Except for Minus One, I don't see too many loners.*

Lenny: *Mingling is only natural at a place like this. And not just because it's cramped. Just keep track of your wallet and don't get too entangled.*

6.1 Mathematical Interlude: Tensor Products

6.1.1 Meet Alice and Bob

Figuring out how systems combine to make bigger systems is a large part of what we do in physics. I hardly need to tell you that an atom is a collection of nucleons and electrons, each of which could be considered a quantum system in its own right.

When talking about composite systems, it's easy to get bogged down in formal language like *System A* and *System B*. Most physicists prefer lighter-weight, informal language instead, and *Alice* and *Bob* have become near-universal substitutes for *A* and *B*. We can think of Alice and Bob as purveyors of composite systems and laboratory setups of every description. Their inventory and expertise are limited only by our imaginations, and they gladly tackle difficult or dangerous assignments like jumping into black holes. They're true geek superheroes!

Let's say that Alice and Bob have provided two systems—Alice's system and Bob's system. Alice's system—whatever it is—is described by a space of states called S_A , and similarly Bob's system is described by a space of states called S_B .

Now let's say that we want to combine the two systems into a single composite system. Before going any further, let's be more specific about the systems we're starting with. For example, Alice's system could be a quantum mechanical coin with two basis states H and T . Of course, a classical coin must be in either one state or the other, but a quantum coin can exist in a superposition:

$$\alpha_H |H\rangle + \alpha_T |T\rangle.$$

You'll notice that I've used an unusual notation for Alice's ket-vectors. This is to distinguish them from Bob's kets. The new notation is intended to discourage us from adding vectors in Alice's space S_A to vectors in Bob's space S_B . Alice's S_A is a two-dimensional vector space—it is defined by the two basis vectors $|H\rangle$ and $|T\rangle$.

Bob's system might also be a coin, but then again it might be something else. Let's assume it's a quantum die. Bob's space of states S_B would then be six-dimensional, with the basis

$$|1\rangle$$
$$|2\rangle$$
$$|3\rangle$$
$$|4\rangle$$
$$|5\rangle$$
$$|6\rangle$$

denoting the six faces of the die. Just like Alice's coin, Bob's die is quantum mechanical, and the six states can be superposed in a similar way.

6.1.2 Representing the Combined System

Now imagine that Bob's and Alice's systems both exist, and form a single composite system. The first question is: How could we construct the state-space—call it S_{AB} —for the combined system? The answer is to form the *tensor product* of S_A and S_B . The notation for this operation is

$$S_{AB} = S_A \otimes S_B.$$

To define S_{AB} , it is enough to specify its basis vectors. The basis vectors are exactly what you might expect. The top

State-Labels for Combined System S_{AB}

		Bob's state-labels					
		1	2	3	4	5	6
Alice's state- labels	<i>H</i>	<i>H1</i>	<i>H2</i>	<i>H3</i>	<i>H4</i>	<i>H5</i>	<i>H6</i>
	<i>T</i>	<i>T1</i>	<i>T2</i>	<i>T3</i>	<i>T4</i>	<i>T5</i>	<i>T6</i>



Alice's
system



Bob's
system

Figure 6.1: The basis states of the composite system S_{AB} , shown as a table. Across the top are the state-labels for Bob's die. The state-labels for Alice's coin are shown on the left. The state-labels for the combined system are the table entries. Each combined state-label shows the state of each of the two subsystems. For example, the state-label $H4$ denotes a state in which Alice's coin shows H and Bob's die shows 4.

half of Fig. 6.1 shows a table whose columns correspond to Bob's six basis vectors and whose rows correspond to Alice's two basis vectors. Each box in the table denotes a basis vector for the S_{AB} system. For example, the box labeled $H4$ represents a state in S_{AB} in which the coin shows Heads and the die shows the number 4. In the combined system, there are twelve basis vectors altogether.

There are various ways to represent these states symbolically. We could represent the $H4$ state using explicit notation, as $|H\rangle \otimes |4\rangle$ or $|H\rangle|4\rangle$. Usually, it's more convenient to use the composite notation $|H4\rangle$. This emphasizes that we're talking about a single state with a two-part label. The left half labels Alice's subsystem, and the right half labels Bob's. The explicit and composite notations both have the same meaning—they refer to the same state.

Once the basis vectors are listed—in this case, twelve of them—we can combine them linearly to form arbitrary superpositions. Thus, the tensor product space in this case is twelve-dimensional. A superposition of two of these basis vectors might look like

$$\alpha_{h3}|H3\rangle + \alpha_{t4}|T4\rangle.$$

In each case, the first half of the state-label describes the state of Alice's coin, and the second half describes the state of Bob's die.

Sometimes, we'll need to refer to an arbitrary basis vector in S_{AB} . To do that, we'll use ket-vectors that look like this,

$$|ab\rangle,$$

or like this,

$$|a'b'\rangle.$$

In this notation, the a or a' (or whatever the left-hand character of the label happens to be) represents one of Alice's states, and the b or b' represents one of Bob's states.

There is one aspect of this notation that is tricky. Even though our S_{AB} state-labels are doubly indexed, ket-vectors like $|ab\rangle$ or $|H3\rangle$ represent a *single state* of the combined system. In other words, we're using a double index to label a single state. This will take some getting used to. Alice's part of the state-label is always on the left and Bob's part is always on the right—keeping Alice and Bob in alphabetical order makes this convention easy to remember.

The rules are the same for more general systems. The only difference is that the two A-states and the six B-states would be replaced by N_A and N_B states respectively, and the tensor product would have dimension

$$N_{AB} = N_A N_B.$$

Systems with three or more components can be represented by tensor products of three or more state spaces, but we won't do that here.

Now that we've described Alice's and Bob's separate spaces S_A and S_B , as well as the combined space S_{AB} , there's still one more bit of notation to set up. Alice has a set of operators, labeled σ , that act on her system. Bob has a similar

set for his system, which we can label τ , so we don't mix them up with Alice's. Alice may have several σ operators, and likewise Bob may have several τ operators. With this framework in hand, we're ready to explore composite systems in greater depth. Later on, in Lecture 7, we'll explain how to work with tensor product operators in component form—expressed as matrices and column vectors.

By now, there should be no doubt in your mind that quantum physics is different from classical physics, right down to its logical roots. In this lecture and the next one, I am going to hit you even harder with this idea. We are going to discuss an aspect of quantum physics that is so different from classical physics that, as of this writing, it has puzzled—and aggravated—physicists and philosophers for almost 80 years. It drove its discoverer, Einstein, to the conclusion that something very deep is missing from quantum mechanics, and physicists have been arguing about it ever since. As Einstein realized, in accepting quantum mechanics, we are buying into a view of reality that is radically different from the classical view.

6.2 Classical Correlation

Before we get to quantum entanglement, let's spend a few minutes on what we might call classical entanglement. In the following experiment, Alice (A) and Bob (B) will get some help from Charlie (C).

Charlie has two coins in his hands—a penny and a dime. He mixes them up and holds them out, one in each hand, to

Alice and Bob, and gives one coin to each of them. No one looks at the coins and no one knows who has which. Then, Alice gets on the shuttle to Alpha Centauri while Bob stays in Palo Alto. Charlie has done his job and doesn't matter anymore (sorry, Charlie).

Before Alice's big trip, Alice and Bob synchronize their clocks—they have done their relativity homework and accounted for time dilation and all that. They agree that Alice will look at her coin just a second or two before Bob looks at his.

Everything proceeds smoothly, and when Alice gets to Alpha Centauri she indeed looks at her coin. Amazingly, the instant she looks at it, she immediately knows exactly what coin Bob will see, even before he looks. Is this crazy? Have Alice and Bob succeeded in breaking relativity's most fundamental rule, which states that information cannot go faster than the speed of light?

Of course not. What would violate relativity would be for Alice's observation *to instantly tell Bob* what to expect. Alice may know what coin Bob will see but she has no way to tell him—not without sending him a real message from Alpha Centauri, and that would take at least the four years required for light to make the trip.

Let's do this experiment many times, either with many Alice-Bob pairs or with the same pair spread out over time. In order to be quantitative, Charlie (he's back now, having accepted our apology) paints a " $\sigma = +1$ " on each penny and a " $\sigma = -1$ " on each dime. If we assume that Charlie really is random in the way he shuffles the coins, then the following facts will emerge:

- On average, both A and B will get as many pennies as dimes. Calling the values of A's observations σ_A and B's observations σ_B , we can express this fact mathematically as

$$\begin{aligned}\langle \sigma_A \rangle &= 0 \\ \langle \sigma_B \rangle &= 0.\end{aligned}\tag{6.1}$$

- If A and B record their observations and then get together back in Palo Alto to compare them, they will find a strong correlation.¹ For each trial, if A observed $\sigma_A = +1$, then B observed $\sigma_B = -1$, and vice versa. In other words, the product $\sigma_A\sigma_B$ always equals -1 :

$$\langle \sigma_A\sigma_B \rangle = -1.$$

Notice that the average of the products (of σ_A and σ_B) is not equal to the product of the averages—Eqs. 6.1 tell us that $\langle \sigma_A \rangle \langle \sigma_B \rangle$ is zero. In symbols,

$$\langle \sigma_A \rangle \langle \sigma_B \rangle \neq \langle \sigma_A\sigma_B \rangle,$$

or

$$\langle \sigma_A\sigma_B \rangle - \langle \sigma_A \rangle \langle \sigma_B \rangle \neq 0.\tag{6.2}$$

¹Actually, it's a *perfect* correlation in this example.

This indicates that Alice's and Bob's observations are *correlated*. In fact, the quantity

$$\langle \sigma_A \sigma_B \rangle - \langle \sigma_A \rangle \langle \sigma_B \rangle$$

is called the statistical correlation between Bob's and Alice's observations. It's called the statistical correlation even if it is zero. When the statistical correlation is nonzero, we say the observations are correlated. The source of this correlation is the fact that originally Alice and Bob were in the same location and Charlie had one of each type of coin. The correlation remained when Alice went to Alpha Centauri simply because the coins didn't change during the trip. There is absolutely nothing strange about this or about Inequality 6.2. It is a very common property of statistical distributions.

Suppose you have a probability distribution $P(a, b)$ for two variables a and b . If the variables are completely uncorrelated, then the probability will factorize:

$$P(a, b) = P_A(a)P_B(b), \quad (6.3)$$

where $P_A(a)$ and $P_B(b)$ are the individual probabilities for a and b . (I added subscripts to the function symbols as a reminder that they could be different functions of their arguments.) It is easy to see that if the probability factorizes in this fashion, then there is no correlation; in other words, the average of the product is the product of the averages.

Exercise 6.1: Prove that if $P(a, b)$ factorizes, then the correlation between a and b is zero.

Let me use an example to illustrate the kind of situation that leads to factorized probabilities. Suppose that instead of a single Charlie, there are two Charlies—Charlie-A and Charlie-B—who have never communicated. Charlie-B mixes up his two coins and gives one to Bob—the other one is discarded.

Charlie-A does exactly the same thing except that he gives a coin to Alice instead. This is the type of situation that leads to factorized product probabilities with no correlation.

In classical physics we use statistics and probability theory when we are ignorant about something that is, in principle, knowable. For example, after mixing up the coins in the first experiment, Charlie could have made a gentle observation (a quick peek) and then let Alice and Bob have their coins. This would have made no difference in the result. In classical mechanics, the probability distribution $P(a, b)$ represents an incomplete specification of the system state. There is more to know—more that could be known—about the system. In classical physics, the use of probability is always associated with an incompleteness of knowledge relative to all that could be known.

A related point is that complete knowledge of a system in classical physics implies complete knowledge of every part of the system. It would not make any sense to say that Charlie knew everything that could be known about the system of two coins but was missing information about the individual coins.

These classical concepts are deeply ingrained in our thinking. They are the foundation of our instinctual understanding of the physical world, and it's very hard to get past

them. But get past them we must, if we are to understand the quantum world.

6.3 Combining Quantum Systems

Charlie's two coins formed a single classical system, composed of two classical subsystems. Quantum mechanics also allows us to combine systems, as we found out in the Mathematical Interlude on tensor products (Section 6.1).

Alice and Bob have kindly agreed to provide a variant of the coin/die system they loaned us for the Interlude on tensor products. Instead of a coin and a die, the new system is built up from two spins, meaning that we'll have a chance to put our knowledge of single spins to work.

As before, we will sometimes use the oddball notation $|a\rangle$ to remind us that Alice's state-vectors are not in the same state-space as Bob's, and that we're not allowed to add them together. On the other hand, recall that each member of an orthonormal basis for S_{AB} is labeled by a pair of vectors, one from S_A and one from S_B . We will make frequent use of the notation $|ab\rangle$ to label a single basis vector of the combined system. These doubly indexed basis vectors *can* be added together, and we'll be doing that a lot.

As we explained in the Interlude, labeling a basis vector with a pair of indices takes some getting used to. You should think of the pair ab as a single index labeling a single state.

Let's look at an example. Consider some linear operator \mathbf{M} acting on the space of states of the composite system. As usual, it can be represented as a matrix. The matrix elements are constructed by sandwiching the operator between

basis vectors. Thus, the matrix elements of \mathbf{M} are expressed as

$$\langle a'b'|\mathbf{M}|ab\rangle = M_{a'b',ab}.$$

Each row of the matrix is labeled with a single index ($a'b'$) of the composite system and each column with (ab).

The vectors $|ab\rangle$ are taken to be orthonormal, which means that their inner products are zero unless both labels match. This does *not* mean that a matches b , but rather that ab matches $a'b'$. We can also express this idea using the Kronecker delta symbol:

$$\langle ab|a'b'\rangle = \delta_{aa'}\delta_{bb'}.$$

The right side is zero unless $a = a'$ and $b = b'$. If the labels do match, the inner product is one.

Now that we have the basis vectors, any linear superposition of them is allowed. Thus, any state in the composite system can be expanded as

$$|\Psi\rangle = \sum_{a,b} \psi(a,b)|ab\rangle.$$

6.4 Two Spins

Returning to our example, let's imagine two spins: Alice's and Bob's. To put it in a context that we can visualize, imagine that the spins are attached to two particles and that the two particles are fixed in space at two nearby but different locations.

Alice and Bob each have their own apparatuses, called \mathcal{A} and \mathcal{B} respectively, that they can use to prepare states and measure spin components. Each can be independently oriented along any axis.

We are going to need names for the two spins. When we only had one spin, we simply called it σ , and it had three components along the x , y , and z axes. Now we have two spins, and the question is how to label them without cluttering the symbols with too many sub- and superscripts. We could call them σ^A and σ^B , and the components, σ_x^A , σ_y^B , and so on. For me, that's just too many subscripts to keep track of, especially on the blackboard. Instead, I'll follow the same convention we used in the Interlude on tensor products. I'll call Alice's spin σ and assign the next letter in the Greek alphabet, τ , to Bob's spin. The full sets of components for Alice's and Bob's spins are

$$\sigma_x, \sigma_y, \sigma_z$$

and

$$\tau_x, \tau_y, \tau_z.$$

According to the principles that we laid out earlier, the space of states for the two-spin system is a tensor product. We can make a table of the four states, just as we did in the Interlude. This time, it's a 2×2 square, comprising four basis states.

Let's work in a basis in which the z components of both spins are specified. The basis vectors are

$$|uu\rangle, |ud\rangle, |du\rangle, |dd\rangle,$$

where the first part of each label represents the state of σ , and the second part represents τ . For example, the first basis vector $|uu\rangle$ represents the state in which both spins are up. The vector $|du\rangle$ is the state in which Alice's spin is down and Bob's spin is up.

6.5 Product States

The simplest type of state for the composite system is called a *product state*. A product state is the result of completely independent preparations by Alice and Bob, in which each uses his or her own apparatus to prepare a spin. Using explicit notation, suppose Alice prepares her spin in state

$$\alpha_u|u\rangle + \alpha_d|d\rangle$$

and Bob prepares his in the state

$$\beta_u|u\rangle + \beta_d|d\rangle.$$

We assume each state is normalized:

$$\begin{aligned}\alpha_u^* \alpha_u + \alpha_d^* \alpha_d &= 1 \\ \beta_u^* \beta_u + \beta_d^* \beta_d &= 1.\end{aligned}\tag{6.4}$$

And in fact these separate normalization equations for each subsystem play a crucial role in defining product states. If they did not hold, we would not have a product state. The product state describing the combined system is

$$|product\ state\rangle = \left\{ \alpha_u|u\rangle + \alpha_d|d\rangle \right\} \otimes \left\{ \beta_u|u\rangle + \beta_d|d\rangle \right\},$$

where the first factor represents Alice's state and the second factor represents Bob's. Expanding the product and switching to composite notation, the right-hand side becomes

$$\alpha_u\beta_u|uu\rangle + \alpha_u\beta_d|ud\rangle + \alpha_d\beta_u|du\rangle + \alpha_d\beta_d|dd\rangle. \quad (6.5)$$

The main feature of a product state is that each subsystem behaves independently of the other. If Bob does an experiment on his own subsystem, the result is exactly the same as it would be if Alice's subsystem did not exist. The same is true for Alice, of course.

Exercise 6.2: Show that if the two normalization conditions of Eqs. 6.4 are satisfied, then the state-vector of Eq. 6.5 is automatically normalized as well. In other words, show that for this product state, normalizing the overall state-vector does not put any additional constraints on the α 's and β 's.

I'll mention here that tensor products and product states are two different things, despite their similar-sounding names.²

²Sometimes, we'll use the term *tensor product space*, or just *product space*, instead of *tensor product*.

A tensor product is a *vector space* for studying composite systems. A product state is a *state-vector*. It's one of the many state-vectors that inhabit a product space. As we will see, most of the state-vectors in the product space are *not* product states.

6.6 Counting Parameters for the Product State

Let's consider the number of parameters it takes to specify such a product state. Each factor requires two complex numbers (α_u and α_d for Alice, β_u and β_d for Bob), which means we need four complex numbers altogether. That's equivalent to eight real parameters. But recall that the normalization conditions in Eqs. 6.4 reduce this by two. Furthermore, the overall phases of each state have no physical significance, so the total number of real parameters is four. That's hardly surprising: it took two parameters to describe the state of a single spin, so two independent spins require four.

6.7 Entangled States

The principles of quantum mechanics allow us to superpose basis vectors in more general ways than just product states. The most general vector in the composite space of states is

$$\psi_{uu}|uu\rangle + \psi_{ud}|ud\rangle + \psi_{du}|du\rangle + \psi_{dd}|dd\rangle,$$

where we have used the subscripted symbols ψ (instead of

α and β) to represent the complex coefficients. Again, we have four complex numbers, but this time we only have one normalization condition,

$$\psi_{uu}^* \psi_{uu} + \psi_{ud}^* \psi_{ud} + \psi_{du}^* \psi_{du} + \psi_{dd}^* \psi_{dd} = 1,$$

and only one overall phase to ignore. The result is that the most general state for a two-spin system has six real parameters. Evidently, the space of states is richer than just those product states that can be prepared independently by Bob and Alice. Something new is going on. The new thing is called *entanglement*.

Entanglement is not an all-or-nothing proposition. Some states are more entangled than others. Here is an example of a maximally entangled state—a state that's as entangled as it can be. It is called the *singlet* state, and it can be written as

$$|\text{sing}\rangle = \frac{1}{\sqrt{2}}(|ud\rangle - |du\rangle).$$

The singlet state cannot be written as a product state. The same is true for the triplet states,

$$\begin{aligned} & \frac{1}{\sqrt{2}}(|ud\rangle + |du\rangle) \\ & \frac{1}{\sqrt{2}}(|uu\rangle + |dd\rangle) \\ & \frac{1}{\sqrt{2}}(|uu\rangle - |dd\rangle), \end{aligned}$$

which are also maximally entangled. The reason for calling them singlet and triplet will be explained later.

Exercise 6.3: Prove that the state $|sing\rangle$ cannot be written as a product state.

What is it about maximally entangled states that is so fascinating? I can sum this up in two statements:

- An entangled state is a complete description of the combined system. No more can be known about it.
- In a maximally entangled state, nothing is known about the individual subsystems.

How can that be? How could we know *as much as can possibly be known* about the Alice-Bob system of two spins, and yet know *nothing* about the individual spins that are its sub-components? That's the mystery of entanglement, and I hope that by the end of this lecture you will understand the rules of the game, even if the deeper nature of entanglement remains a paradox.

6.8 Alice and Bob's Observables

So far, we've discussed the space of states of the Alice-Bob two-spin system, but not its observables. Some of these observables are obvious, even if their mathematical representation is not. In particular, using their apparatuses \mathcal{A} and \mathcal{B} , Alice and Bob can measure the components of their spins:

$$\sigma_x, \sigma_y, \sigma_z$$

and

$$\tau_x, \tau_y, \tau_z.$$

How are these observables represented as Hermitian operators in the composite space of states? The answer is simple. Bob's operators act on Bob's spin states exactly as they would if Alice had never shown up. The same goes for Alice. Let's review how the spin operators act on the states of a single spin. First, let's look at Alice's spin:

$$\begin{aligned} \sigma_z|u\rangle &= |u\rangle \\ \sigma_z|d\rangle &= -|d\rangle \\ \sigma_x|u\rangle &= |d\rangle \\ \sigma_x|d\rangle &= |u\rangle \\ \sigma_y|u\rangle &= i|d\rangle \\ \sigma_y|d\rangle &= -i|u\rangle. \end{aligned} \tag{6.6}$$

Of course, Bob's setup is identical to Alice's, so we can write a parallel set of equations showing how the components of τ act on Bob's states:

$$\begin{aligned}
\tau_z|u\rangle &= |u\rangle \\
\tau_z|d\rangle &= -|d\rangle \\
\tau_x|u\rangle &= |d\rangle \\
\tau_x|d\rangle &= |u\rangle \\
\tau_y|u\rangle &= i|d\rangle \\
\tau_y|d\rangle &= -i|u\rangle.
\end{aligned} \tag{6.7}$$

Now let's consider how the operators should be defined when acting on the tensor product states, $|uu\rangle$, $|ud\rangle$, $|du\rangle$, and $|dd\rangle$. The answer is that when σ acts, it just ignores Bob's half of the state label. There are many possible combinations of operators and states, but I will pick a few at random. You can fill in the others, or look them up in the appendix. Starting with Alice's operators, we find that

$$\begin{aligned}
\sigma_z|uu\rangle &= |uu\rangle \\
\sigma_z|du\rangle &= -|du\rangle \\
\sigma_x|ud\rangle &= |dd\rangle \\
\sigma_x|dd\rangle &= |ud\rangle
\end{aligned}$$

$$\begin{aligned}
\sigma_y|uu\rangle &= i|du\rangle \\
\sigma_y|du\rangle &= -i|uu\rangle \\
\tau_z|uu\rangle &= |uu\rangle \\
\tau_z|du\rangle &= |du\rangle \\
\tau_x|ud\rangle &= |uu\rangle \\
\tau_x|du\rangle &= |dd\rangle \\
\tau_y|uu\rangle &= i|ud\rangle \\
\tau_y|dd\rangle &= -i|du\rangle.
\end{aligned} \tag{6.8}$$

Again, the rule is that Alice's spin components act only on the Alice half of the composite system. The Bob half is a passive spectator that does not participate. In terms of symbols, when σ_x , σ_y , or σ_z acts, Bob's half of the spin state does not change. And when Bob's τ spin operators act, Alice's half is similarly passive.

We are being a little loose with our notation. The vectors of a tensor product space are *new* vectors, built up from the vectors of two smaller spaces. Technically, the same is true for the operators. If we were being pedantic, we would insist on writing the tensor product versions of σ_z and τ_x as $\sigma_z \otimes I$ and $I \otimes \tau_x$, respectively, where I is the identity operator. In fact, we can highlight two important properties of tensor

product operators by rewriting the equation

$$\sigma_z|du\rangle = -|du\rangle \tag{6.9}$$

as

$$\begin{aligned} (\sigma_z \otimes I) (|d\rangle \otimes |u\rangle) &= (\sigma_z|d\rangle \otimes I|u\rangle) \\ &= (-|d\rangle \otimes |u\rangle). \end{aligned} \tag{6.10}$$

This notation is cumbersome, and we'll usually stick to the simpler language of Eq. 6.9. However, the language of Eq. 6.10 makes two things clear:

1. A composite operator $\sigma_z \otimes I$ is operating on a composite vector $|d\rangle \otimes |u\rangle$ to produce a new composite vector $-|d\rangle \otimes |u\rangle$.
2. Alice's half (the left half) of the composite operator only affects her half of the composite vector. Likewise, Bob's half of the operator only affects his half of the vector.

We'll have more to say about composite operators in the next section. Furthermore, in Lecture 7, the language of Eq. 6.10 will help us see how to work with tensor products in component form.

Exercise 6.4: Use the matrix forms of σ_z , σ_x , and σ_y and the column vectors for $|u\rangle$ and $|d\rangle$ to verify Eqs. 6.6. Then, use Eqs. 6.6 and 6.7 to write the equations that were left out of Eqs. 6.8. Use the appendix to check your answers.

Exercise 6.5: Prove the following theorem:

When any one of Alice's or Bob's spin operators acts on a product state, the result is still a product state.

Show that in a product state, the expectation value of any component of $\vec{\sigma}$ or $\vec{\tau}$ is exactly the same as it would be in the individual single-spin states.

This last exercise proves something important about product states. In a product state, every prediction about Bob's half of the system is exactly the same as it would have been in the corresponding single-spin theory. The same goes for Alice.

An example of this property of product states involves what I called the *Spin-Polarization Principle* in Lecture 3. A useful way to state that principle is:

For any state of a single spin, there is some direction for which the spin is +1.

As I explained, this means that the expectation values of the components satisfy the equation

$$\langle\sigma_x\rangle^2 + \langle\sigma_y\rangle^2 + \langle\sigma_z\rangle^2 = 1, \quad (6.11)$$

which tells us that not all the expectation values can be zero. This fact continues to hold for all product states. However,

it does *not* hold for the entangled state $|sing\rangle$. In fact, for the $|sing\rangle$ state the right-hand side of Eq. 6.11 becomes zero, as we'll show next.

Recall that the entangled state $|sing\rangle$ is defined as

$$|sing\rangle = \frac{1}{\sqrt{2}}(|ud\rangle - |du\rangle).$$

Let's look at the expectation values of σ in this state. We have all the machinery we need to compute them. First, let's consider $\langle\sigma_z\rangle$:

$$\begin{aligned}\langle\sigma_z\rangle &= \langle sing|\sigma_z|sing\rangle \\ &= \langle sing|\sigma_z\frac{1}{\sqrt{2}}(|ud\rangle - |du\rangle).\end{aligned}$$

Here is where Eqs. 6.8 come in (along with Exercise 6.4, which completes this set of equations!). They tell us how σ_z acts on each basis vector. The result is

$$\langle sing|\sigma_z|sing\rangle = \langle sing|\frac{1}{\sqrt{2}}(|ud\rangle + |du\rangle)$$

or

$$\langle\sigma_z\rangle = \frac{1}{2}\left(\langle ud| - \langle du|\right)\left(|ud\rangle + |du\rangle\right).$$

A quick inspection shows that this is equal to zero. Next, let's consider $\langle\sigma_x\rangle$:

$$\begin{aligned}\langle \sigma_x \rangle &= \langle \text{sing} | \sigma_x | \text{sing} \rangle \\ &= \langle \text{sing} | \sigma_x \frac{1}{\sqrt{2}} (|ud\rangle - |du\rangle) \end{aligned}$$

or

$$\langle \sigma_x \rangle = \frac{1}{2} (\langle ud| - \langle du|) (|dd\rangle - |uu\rangle).$$

Again, this equation gives us zero. Finally, let's look at $\langle \sigma_y \rangle$:

$$\begin{aligned}\langle \sigma_y \rangle &= \langle \text{sing} | \sigma_y | \text{sing} \rangle \\ &= \frac{1}{2} (\langle ud| - \langle du|) (i|dd\rangle + i|uu\rangle). \end{aligned}$$

As you may have guessed, we are left with zero once more. Thus, we have shown that for the state $|\text{sing}\rangle$,

$$\langle \sigma_z \rangle = \langle \sigma_x \rangle = \langle \sigma_y \rangle = 0,$$

and indeed all expectation values of σ are zero. Needless to say, the same is true for the expectation values of τ . Clearly, $|\text{sing}\rangle$ is very different from a product state. What does all this say about the measurements we can make?

If the expectation value of a component of σ is zero, it means that the experimental outcome is equally likely to be

+1 or -1 . In other words, the outcome is completely uncertain. Even though we know the exact state-vector, $|sing\rangle$, we know nothing at all about the outcome of any measurement of any component of either spin.

Perhaps this means that the state $|sing\rangle$ is somehow incomplete—that there are details of the system that we were sloppy about and didn't measure. After all, earlier we saw a perfectly classical example in which Alice and Bob knew nothing about their coins until they actually looked at them. How is the quantum version different?

In our “classical entanglement” example involving Alice, Bob, and Charlie, it is perfectly clear that there was more to know. Charlie could have sneaked a peek at the coins without changing anything, because classical measurements can be arbitrarily gentle.

Might there be so-called *hidden variables* in the quantum system? The answer is that according to the rules of quantum mechanics, there is nothing to know beyond what is encoded in the state-vector—in the present case, $|sing\rangle$. The state-vector is as complete a description of a system as it is possible to make. So it seems that in quantum mechanics, we can know everything about a composite system—everything there is to know, anyway—and still know nothing about its constituent parts. This is the true weirdness of entanglement, which so disturbed Einstein.

6.9 Composite Observables

Let's imagine a quantum mechanical Alice-Bob-Charlie setup. Charlie's role is to prepare two spins in the entangled state

sing). Then, without looking at the spins (remember, quantum measurements are *not* gentle), he gives one spin to Alice and one to Bob. Although Alice and Bob know exactly what state the combined system is in, they can predict nothing about the outcome of their individual measurements.

But surely knowing the exact state of the composite system must tell them *something*, even if the state is highly entangled. And in fact it does. However, to understand what it tells them, we have to consider a wider family of observables than the ones that Alice and Bob can measure separately, each using *only* his or her own detector. As it turns out, there are observables that can only be measured by using both detectors. The results of such experiments can only be known to Alice or Bob if they come together and compare notes.

The first question is whether Alice and Bob can simultaneously measure their own observables. We have seen that there are quantities that cannot be simultaneously measured. In particular, two observables that do not commute cannot both be measured without the measurements interfering with each other. But for Alice and Bob, it is easy to see that every component of σ commutes with every component of τ . This is a general fact about tensor products. The operators that act on the two separate factors commute with one another. Therefore, Alice can make any measurement on her spin and Bob can make any measurement on his, without either interfering with the other's experiment.

Let's suppose Alice measures σ_z and Bob measures τ_z , and then they multiply the results. In other words, they conspire to measure the product $\tau_z\sigma_z$.

The product $\tau_z\sigma_z$ is an observable that is mathematically represented by first applying σ_z to a ket and then subsequently applying τ_z . Keep in mind that these are just the mathematical operations that define a new operator: they are different from the act of performing a physical measurement. You don't need an apparatus to multiply two operators; you just need a pencil and paper. Let's see what happens if we apply the product $\tau_z\sigma_z$ to the state $|sing\rangle$:

$$\tau_z\sigma_z\frac{1}{\sqrt{2}}(|ud\rangle - |du\rangle).$$

First, using the table in Eqs. 6.8, apply σ_z :

$$\tau_z\sigma_z\frac{1}{\sqrt{2}}(|ud\rangle - |du\rangle) = \tau_z\frac{1}{\sqrt{2}}(|ud\rangle + |du\rangle).$$

Now, apply τ_z to get

$$\tau_z\sigma_z\frac{1}{\sqrt{2}}(|ud\rangle - |du\rangle) = \frac{1}{\sqrt{2}}(-|ud\rangle + |du\rangle).$$

Notice that the end result is just to change the sign of $|sing\rangle$:

$$\tau_z\sigma_z|sing\rangle = -|sing\rangle.$$

Evidently, $|sing\rangle$ is an eigenvector of the observable $\tau_z\sigma_z$ with eigenvalue -1 . Let's examine the significance of this result. Alice measures σ_z and Bob measures τ_z ; when they come together and compare results, they find they've measured opposite values. Sometimes, Bob measures $+1$ and

Alice measures -1 . Other times, Alice measures $+1$ and Bob measures -1 . The product of the two measurements is always -1 .

There should be nothing surprising in this result. The state-vector $|sing\rangle$ is a superposition of two vectors, $|ud\rangle$ and $|du\rangle$, both of which comprise two spins with opposite z components. The situation is altogether similar to the classical example involving Charlie and his two coins.

But now we come to something that has no classical analog. Suppose that instead of measuring the z components of their spins, Alice and Bob measure the x components. To find out how their outcomes are correlated, we must study the observable $\tau_x\sigma_x$.

Let's act on $|sing\rangle$ with this product. Here are the steps:

$$\begin{aligned}\tau_x\sigma_x|sing\rangle &= \tau_x\sigma_x\frac{1}{\sqrt{2}}(|ud\rangle - |du\rangle) \\ &= \tau_x\frac{1}{\sqrt{2}}(|dd\rangle - |uu\rangle) \\ &= \frac{1}{\sqrt{2}}(|du\rangle - |ud\rangle)\end{aligned}$$

or, more simply,

$$\tau_x\sigma_x|sing\rangle = -|sing\rangle.$$

Now this is a bit surprising: $|sing\rangle$ is also an eigenvector of $\tau_x\sigma_x$ with eigenvalue -1 . It is far less obvious from just looking at $|sing\rangle$ that the x components of the two spins

are always opposite. Nevertheless, every time Alice and Bob measure them, they find that σ_x and τ_x have opposite values. At this point, you will probably not be surprised to learn that the same thing is true for the y components.

Exercise 6.6: Assume Charlie has prepared the two spins in the singlet state. This time, Bob measures τ_y and Alice measures σ_x . What is the expectation value of $\sigma_x\tau_y$?

What does this say about the correlation between the two measurements?

Exercise 6.7: Next, Charlie prepares the spins in a different state, called $|T_1\rangle$, where

$$|T_1\rangle = \frac{1}{\sqrt{2}}(|ud\rangle + |du\rangle).$$

In these examples, T stands for *triplet*. These triplet states are completely different from the states in the coin and die examples. What are the expectation values of the operators $\sigma_z\tau_z$, $\sigma_x\tau_x$, and $\sigma_y\tau_y$?

What a difference a sign can make!

Exercise 6.8: Do the same for the other two entangled triplet states,

$$|T_2\rangle = \frac{1}{\sqrt{2}}(|uu\rangle + |dd\rangle)$$

$$|T_3\rangle = \frac{1}{\sqrt{2}}(|uu\rangle - |dd\rangle),$$

and interpret.

Finally, let's consider one more observable. This one cannot be measured by Alice and Bob making separate measurements with their individual apparatuses, even if they come together and compare notes. Nevertheless, quantum mechanics insists that some kind of apparatus can be built to measure the observable.

The observable I am referring to can be thought of as the ordinary dot product of the vector-operators $\vec{\sigma}$ and $\vec{\tau}$:

$$\vec{\sigma} \cdot \vec{\tau} = \sigma_x \tau_x + \sigma_y \tau_y + \sigma_z \tau_z.$$

One might think that a value for this observable can be found if Bob measures all components of τ , while Alice measures all components of σ ; then they could multiply the components and add them up. The problem is that Bob cannot simultaneously measure the individual components of τ , because they don't commute. Likewise, Alice cannot measure more than one component of σ at a time. To measure $\vec{\sigma} \cdot \vec{\tau}$, a new kind of apparatus must be built, one that measures $\vec{\sigma} \cdot \vec{\tau}$ without measuring any individual component. It's far from obvious how that could be done. Here is a concrete example of how such a measurement could be carried out: Some atoms have spins that are described in the same way as electron spins. When two of these atoms are close to each other—for example, two neighboring atoms in a crystal lattice—the Hamiltonian will depend on the spins. In some situations, the neighboring spins' Hamiltonian is proportional to $\vec{\sigma} \cdot \vec{\tau}$. If that happens to be the case, then measuring $\vec{\sigma} \cdot \vec{\tau}$ is equivalent to measuring the energy of the atomic pair. Measuring this energy is a single measurement of the composite opera-

tor and does not entail measuring the individual components of either spin.

Exercise 6.9: Prove that the four vectors $|sing\rangle$, $|T_1\rangle$, $|T_2\rangle$, and $|T_3\rangle$ are eigenvectors of $\vec{\sigma} \cdot \vec{\tau}$. What are their eigenvalues?

Take a look at your results from this last exercise. Do you see why one of these state-vectors is called the singlet, while the other three are called triplets? The reason is that if you look at their relation to the operator $\vec{\sigma} \cdot \vec{\tau}$, the singlet is an eigenvector with one eigenvalue, and the triplets are all eigenvectors with a different *degenerate* eigenvalue.

Here is a good exercise that combines the concept of entanglement with the concepts of time and change from Lecture 4. Use it to review the ideas of unitary time evolution and the meaning of the Hamiltonian.

Exercise 6.10: A system of two spins has the Hamiltonian

$$\mathbf{H} = \frac{\omega}{2} \vec{\sigma} \cdot \vec{\tau}.$$

What are the possible energies of the system, and what are the eigenvectors of the Hamiltonian?

Suppose the system starts in the state $|uu\rangle$. What is the state at any later time? Answer the same question for initial states of $|ud\rangle$, $|du\rangle$, and $|dd\rangle$.

Lecture 7

More on Entanglement

Hilbert's Place, summer 1935:

Two scruffy regulars come through the swinging doors, in the midst of an intense conversation. The one with the wild grayish hair and frayed sweater says, "No, I will not accept your theory unless you can tell me what the elements of physical reality are."

The other one looks around, throws up his hands in obvious frustration, and says to Art and Lenny, "There he goes again. Elements of physical reality, EPRs, EPRs, that's all he ever thinks about. Albert, stop being obsessive and just accept the facts."

"Never! I cannot accept that one can know everything there is to know about a thing, and still know nothing about its parts. That's utter nonsense, Niels."

"Sorry, Albert. That's just the way it is. Here, let me buy you a beer."

In this lecture, we will look at entanglement in greater depth. To do that, we'll need some additional mathematical tools. First, we'll find out how to work with tensor products in component form. Then, we'll learn about a new operator called the *density matrix*. These tools are not inherently hard to master, but they do require some patience and a fair amount of index wrangling.

7.1 Mathematical Interlude: Tensor Products in Component Form

In Lecture 6, we explained how to form the tensor product of two vector spaces using the abstract notation of bras, kets, and operator symbols like σ_z . How does that translate into columns, rows, and matrices?

Building tensor products from matrices and column vectors is not hard. The rules are straightforward, as we'll see below. The tricky part is understanding why these rules work—why they allow us to build matrices and column vectors *that have the properties we want*. We'll tackle the issue in two different ways. First, we'll build composite operators using the tried-and-true method we developed in Lecture 3. Then we'll show you how to build composite operators directly from their component operators.

7.1.1 Building Tensor Product Matrices from Basic Principles

Back in Lecture 3, we showed you how to write *any* observable \mathbf{M} in matrix form, relative to a specific basis. Take a moment to review Eqs. 3.1 through 3.4. In that section, we calculated the numerical values m_{jk} of \mathbf{M} 's matrix elements with the expression

$$m_{jk} = \langle j | \mathbf{M} | k \rangle, \quad (7.1)$$

where $|j\rangle$ and $|k\rangle$ represent the basis vectors. Each $|j\rangle, |k\rangle$ combination generates a different matrix element.¹

Our plan is to apply this formula to some tensor product operators and see what we get. Because of our double-indexing convention for tensor product basis vectors, the “sandwiches” in these equations will look a little different from the ones in Eq. 7.1. On each end of the sandwich, we will cycle through the basis vectors $|uu\rangle, |ud\rangle, |du\rangle$, and $|dd\rangle$.² To keep things simple, we'll use the operator $\sigma_z \otimes I$ as an example, where I is the identity operator. As we have seen, $\sigma_z \otimes I$ acts on Alice's half of the state-vector with σ_z , and does *absolutely nothing* to Bob's half. Because we are working in a four-dimensional vector space, the resulting matrix

¹In Lecture 3, we happened to write the index j on the left side of \mathbf{M} , and k on the right, the opposite of what we're doing here. Because j and k are index variables, this makes no difference as long as we maintain consistency within a group of equations.

²Of course, we could have used a different set of basis vectors, such as $|rr\rangle, |rl\rangle$, etc. Doing so would result in a different set of matrix elements.

will be 4×4 . Omitting multiple \otimes symbols to avoid visual clutter, we can write the matrix like this:

$$\sigma_z \otimes I = \begin{pmatrix} \langle uu|\sigma_z I|uu\rangle & \langle uu|\sigma_z I|ud\rangle & \langle uu|\sigma_z I|du\rangle & \langle uu|\sigma_z I|dd\rangle \\ \langle ud|\sigma_z I|uu\rangle & \langle ud|\sigma_z I|ud\rangle & \langle ud|\sigma_z I|du\rangle & \langle ud|\sigma_z I|dd\rangle \\ \langle du|\sigma_z I|uu\rangle & \langle du|\sigma_z I|ud\rangle & \langle du|\sigma_z I|du\rangle & \langle du|\sigma_z I|dd\rangle \\ \langle dd|\sigma_z I|uu\rangle & \langle dd|\sigma_z I|ud\rangle & \langle dd|\sigma_z I|du\rangle & \langle dd|\sigma_z I|dd\rangle \end{pmatrix}. \quad (7.2)$$

To evaluate these matrix elements, we could allow σ_z and I to operate either to the left or to the right. Let's assume σ_z operates to the left and I operates to the right. Since I does nothing, all we care about is what σ_z does to the bra vector on its left. And within that bra vector, σ_z only acts on the leftmost (that is, Alice's) state-label. Using the rules we've already worked out (see Eqs. 6.6 and 6.7), we can carry out all of these σ_z operations to obtain a matrix of inner products:

$$\sigma_z \otimes I = \begin{pmatrix} \langle uu|uu\rangle & \langle uu|ud\rangle & \langle uu|du\rangle & \langle uu|dd\rangle \\ \langle ud|uu\rangle & \langle ud|ud\rangle & \langle ud|du\rangle & \langle ud|dd\rangle \\ -\langle du|uu\rangle & -\langle du|ud\rangle & -\langle du|du\rangle & -\langle du|dd\rangle \\ -\langle dd|uu\rangle & -\langle dd|ud\rangle & -\langle dd|du\rangle & -\langle dd|dd\rangle \end{pmatrix}. \quad (7.3)$$

Because these eigenvectors are orthonormal, the matrix reduces to

$$\sigma_z \otimes I = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}. \quad (7.4)$$

How do we write the eigenvectors $|uu\rangle$, $|ud\rangle$, $|du\rangle$, and $|dd\rangle$ as column vectors? For now, I'll just tell you that we'll represent $|uu\rangle$ and $|du\rangle$ as

$$|uu\rangle = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \quad |du\rangle = \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix}. \quad (7.5)$$

Let's see what happens when $\sigma_z \otimes I$ operates on these column vectors. Applying the matrix to $|uu\rangle$ results in

$$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}.$$

In other words,

$$(\sigma_z \otimes I)|uu\rangle = |uu\rangle,$$

just as we expect. What if we apply the same matrix to the column vector $|du\rangle$ in Eqs. 7.5? Carrying out the matrix multiplication results in $-|du\rangle$, just as it should.

7.1.2 Building Tensor Product Matrices from Component Matrices

The above method for calculating matrix elements is very general—it works for *all* observables. If we need to construct the tensor product of two operators, and we *already know* the matrix elements of the building blocks, we can combine them directly. Here is the rule for combining 2×2 matrices to form 4×4 matrices:

$$A \otimes B = \begin{pmatrix} A_{11}B & A_{12}B \\ A_{21}B & A_{22}B \end{pmatrix} \quad (7.6)$$

or

$$A \otimes B = \begin{pmatrix} A_{11}B_{11} & A_{11}B_{12} & A_{12}B_{11} & A_{12}B_{12} \\ A_{11}B_{21} & A_{11}B_{22} & A_{12}B_{21} & A_{12}B_{22} \\ A_{21}B_{11} & A_{21}B_{12} & A_{22}B_{11} & A_{22}B_{12} \\ A_{21}B_{21} & A_{21}B_{22} & A_{22}B_{21} & A_{22}B_{22} \end{pmatrix}. \quad (7.7)$$

The same pattern works for matrices of any size. This kind of matrix multiplication is sometimes called the *Kronecker product*, a term that only applies to matrices—it's the matrix version of the tensor product. The Kronecker product of two 2×2 matrices is a 4×4 matrix, and the pattern is similar for matrices of arbitrary size. In general, the Kronecker product of an $m \times n$ matrix and a $p \times q$ matrix is an $mp \times nq$ matrix.

All of this applies perfectly well to column and row vectors, which are just specialized matrices. The tensor product of two 2×1 column vectors is a 4×1 column vector. If a and b are 2×1 column vectors, their tensor product looks like this:

$$\begin{pmatrix} a_{11} \\ a_{21} \end{pmatrix} \otimes \begin{pmatrix} b_{11} \\ b_{21} \end{pmatrix} = \begin{pmatrix} a_{11}b_{11} \\ a_{11}b_{21} \\ a_{21}b_{11} \\ a_{21}b_{21} \end{pmatrix}. \quad (7.8)$$

Let's see how this works out for Alice and Bob. First, we'll construct the four tensor product basis vectors, using $|u\rangle$ and $|d\rangle$ as building blocks. Recall Eqs. 2.11 and 2.12 from Lecture 2,

$$|u\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$$

$$|d\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

If we plug the appropriate combinations of $|u\rangle$ and $|d\rangle$ into Eq. 7.8, our four 4×1 column vectors are

$$|uu\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}$$

$$|ud\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \otimes \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix}$$

$$|du\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix}$$

$$|dd\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \otimes \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}. \quad (7.9)$$

Next, we'll use the rule from Eq. 7.7 to combine the operators σ_z and τ_x . Using Eqs. 3.20 to define matrices σ_z and τ_x , this rule gives the tensor product matrix

$$\sigma_z \otimes \tau_x = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \otimes \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 \\ 0 & 0 & -1 & 0 \end{pmatrix}.$$

Let's compare this result with the product of σ_x and τ_z ,

$$\sigma_x \otimes \tau_z = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \otimes \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} = \begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \\ 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \end{pmatrix}.$$

Notice that $\sigma_x \otimes \tau_z$ is not the same as $\sigma_z \otimes \tau_x$. That is natural, because they represent different observables.

So far, so good. But next, we'll see something a little more interesting. With the help of a few exercises, we'll try to convince you that the Kronecker product really *is* the tensor product for matrices—in other words, that Alice's half of the matrix only affects her half of the column vector, and likewise for Bob. This is tricky because of the way the Kronecker product mixes up the elements of its building blocks.

As an example, let's look at how $\sigma_z \otimes \tau_x$ acts on $|ud\rangle$. Translating the abstract symbols into components, we can write

$$(\sigma_z \otimes \tau_x)|ud\rangle = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 \\ 0 & 0 & -1 & 0 \end{pmatrix} \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}.$$

But the column vector on the right-hand side corresponds to $|uu\rangle$ in Eqs. 7.9. Translated back into abstract notation, this becomes

$$(\sigma_z \otimes \tau_x)|ud\rangle = |uu\rangle.$$

This is exactly what we want—a matrix representation of our abstract operators and state-vectors that replicates their known behavior.

The following exercise will help crystallize the idea that the σ -half of $\sigma \otimes \tau$ only affects Alice's half of the state-vector, and that the τ -half only affects Bob's. The one after that provides some practice working out the matrix elements of an operator, assuming that we already know what the operator does to each basis vector.

Exercise 7.1: Write the tensor product $I \otimes \tau_x$ as a matrix, and apply that matrix to each of the $|uu\rangle$, $|ud\rangle$, $|du\rangle$, and $|dd\rangle$ column vectors. Show that Alice's half of the state-vector is unchanged in each case. Recall that I is the 2×2 unit matrix.

Exercise 7.2: Calculate the matrix elements of $\sigma_z \otimes \tau_x$ by forming inner products as we did in Eq. 7.2.

The third exercise is a bit tedious, but it really nails things down. Consider the equation

$$(A \otimes B) (a \otimes b) = (Aa \otimes Bb). \quad (7.10)$$

As in Eqs. 7.7 and 7.8, A and B represent 2×2 matrices (or operators), and a and b represent 2×1 column vectors. The exercise asks you to expand the equation into components and show that the left side matches the right side.

Exercise 7.3:

- a) Rewrite Eq. 7.10 in component form, replacing the symbols A , B , a , and b with the matrices and column vectors from Eqs. 7.7 and 7.8.
- b) Perform the matrix multiplications Aa and Bb on the right-hand side. Verify that each result is a 4×1 matrix.
- c) Expand all three Kronecker products.
- d) Verify the row and column sizes of each Kronecker product:
 - $A \otimes B$: 4×4
 - $a \otimes b$: 4×1
 - $Aa \otimes Bb$: 4×4
- e) Perform the matrix multiplication on the left-hand side, resulting in a 4×1 column vector. Each row should be the sum of four separate terms.
- f) Finally, verify that the resulting column vectors on the left and right sides are identical.

7.2 Mathematical Interlude: Outer Products

Given a bra $\langle\phi|$ and a ket $|\psi\rangle$, we can form the inner product $\langle\phi|\psi\rangle$. As we've seen, the inner product is a complex number. However, there is another kind of product called the outer product, written

$$|\psi\rangle\langle\phi|.$$

The outer product is not a number; it is a linear operator. Let's consider what happens when $|\psi\rangle\langle\phi|$ acts on another ket $|A\rangle$:

$$|\psi\rangle\langle\phi| |A\rangle.$$

In these examples, we're using spacing instead of parentheses to show the grouping of operations. Remember that all operations with bras, kets, and linear operators are associative, which means we're allowed to group them any way we like, as long as we keep the same ordering from left to right.³ The action of the outer product operator is very simple and can be defined as

$$|\psi\rangle\langle\phi| |A\rangle \equiv |\psi\rangle \langle\phi|A\rangle.$$

³Sometimes we can change left-to-right ordering as well, but that requires more care.

In other words, we take the inner product of $\langle\phi|$ with $|A\rangle$ (the result is a complex number) and multiply it by the ket $|\psi\rangle$. The bra-ket notation is so efficient that it practically forces the definition on us. That was the genius of Paul Dirac. It's easy to prove that the outer product can also act on bras:

$$\langle B| |\psi\rangle\langle\phi| \equiv \langle B|\psi\rangle \langle\phi|.$$

A special case is the outer product of a ket with its corresponding bra, $|\psi\rangle\langle\psi|$. Assuming that $|\psi\rangle$ is normalized, this operator is called a *projection operator*. Here is how it acts:

$$|\psi\rangle\langle\psi| |A\rangle = |\psi\rangle \langle\psi|A\rangle$$

Note that the result is always proportional to $|\psi\rangle$. A projection operator can be said to project a vector onto the direction defined by $|\psi\rangle$. Here are some properties of projection operators that you can easily prove (remember that $|\psi\rangle$ is normalized to 1):

- Projection operators are Hermitian.
- The vector $|\psi\rangle$ is an eigenvector of its projection operator with eigenvalue 1:

$$|\psi\rangle\langle\psi| |\psi\rangle = |\psi\rangle$$

- Any vector orthogonal to $|\psi\rangle$ is an eigenvector with eigenvalue zero. Thus, the eigenvalues of $|\psi\rangle\langle\psi|$ are all either 0 or 1, and there is only one eigenvector with unit eigenvalue. That eigenvector is $|\psi\rangle$ itself.

- The square of a projection operator is the same as the projection operator itself:

$$|\psi\rangle\langle\psi|^2 = |\psi\rangle\langle\psi|.$$

- The *trace* of an operator (or any square matrix) is defined as the sum of its diagonal elements. Using the symbol Tr for trace, we can define the trace of an operator \mathbf{L} to be

$$Tr \mathbf{L} = \sum_i \langle i | \mathbf{L} | i \rangle,$$

which is just the sum of \mathbf{L} 's diagonal matrix elements.

The trace of a projection operator is 1. This follows from the fact that the trace of a Hermitian operator is the sum of its eigenvalues.⁴

- If we add all the projection operators for a basis system, we obtain the identity operator:

$$\sum_i |i\rangle\langle i| = I. \quad (7.11)$$

Finally, here is a very important theorem about projection operators and expectation values. The expectation value of

⁴A Hermitian matrix \mathbf{M} can be diagonalized by a transformation $\mathbf{P}^\dagger \mathbf{M} \mathbf{P}$, where \mathbf{P} is a unitary matrix whose columns are the normalized eigenvectors of \mathbf{M} . The trace of \mathbf{M} is invariant under this transformation. We have not proved this well-known result.

any observable \mathbf{L} in state $|\psi\rangle$ is given by

$$\langle\psi|\mathbf{L}|\psi\rangle = \text{Tr} |\psi\rangle\langle\psi| \mathbf{L}. \quad (7.12)$$

Here are the steps to prove it. Pick any basis $|i\rangle$. Then, using the definition of *trace*, write

$$\text{Tr} |\psi\rangle\langle\psi| \mathbf{L} = \sum_i \langle i|\psi\rangle\langle\psi|\mathbf{L}|i\rangle.$$

The two factors in the summation are just numbers, so we can reverse their ordering,

$$\text{Tr} |\psi\rangle\langle\psi| \mathbf{L} = \sum_i \langle\psi|\mathbf{L}|i\rangle\langle i|\psi\rangle.$$

Carrying out the sum and using $\sum |i\rangle\langle i| = I$, we get

$$\text{Tr} |\psi\rangle\langle\psi| \mathbf{L} = \langle\psi|\mathbf{L}|\psi\rangle.$$

The right side is just the expectation value of \mathbf{L} .

7.3 Density Matrices: A New Tool

Up to now, we have learned how to make predictions about a system when we know the system's exact quantum state. But more often than not, we don't have complete knowledge of the state. For example, suppose Alice has prepared a spin using an apparatus oriented along some axis. She gives the spin to Bob but doesn't tell him the axis along which the apparatus was oriented. Perhaps she gives him some partial

information, such as the fact that the axis was either along the z axis or the x axis, but she refuses to tell him more than that. What does Bob do? How does he use this information to make predictions?

Bob reasons as follows: If Alice prepared the spin in the state $|\psi\rangle$, then the expectation value of any observable \mathbf{L} is

$$\text{Tr } |\psi\rangle\langle\psi|\mathbf{L} = \langle\psi|\mathbf{L}|\psi\rangle.$$

On the other hand, if Alice prepared the spin in state $|\phi\rangle$, then the expectation value of \mathbf{L} is

$$\text{Tr } |\phi\rangle\langle\phi|\mathbf{L} = \langle\phi|\mathbf{L}|\phi\rangle.$$

What if there is a 50 percent probability that she prepared $|\psi\rangle$ and a 50 percent probability that she prepared $|\phi\rangle$? Obviously, the expectation value is

$$\langle\mathbf{L}\rangle = \frac{1}{2}\text{Tr } |\psi\rangle\langle\psi|\mathbf{L} + \frac{1}{2}\text{Tr } |\phi\rangle\langle\phi|\mathbf{L}.$$

All we are doing is averaging over Bob's ignorance of the state prepared by Alice.

But now we can combine the terms into a single expression by defining a density matrix ρ that encodes Bob's knowledge. In this case the density matrix is half the projection operator onto $|\phi\rangle$ plus half the projection operator onto $|\psi\rangle$,

$$\rho = \frac{1}{2}|\psi\rangle\langle\psi| + \frac{1}{2}|\phi\rangle\langle\phi|.$$

We've now packaged all of Bob's knowledge of the system into a single operator ρ . At this point, the rule to compute expectation values becomes very simple:

$$\langle L \rangle = \text{Tr } \rho \mathbf{L}. \quad (7.13)$$

We can generalize this. Suppose that Alice tells Bob that she has prepared one of several states—call them $|\phi_1\rangle$, $|\phi_2\rangle$, $|\phi_3\rangle$, and so on. Moreover, she specifies probabilities P_1 , P_2 , P_3, \dots for each of these states. Bob can still package all his knowledge into a density matrix:

$$\rho = P_1|\phi_1\rangle\langle\phi_1| + P_2|\phi_2\rangle\langle\phi_2| + P_3|\phi_3\rangle\langle\phi_3| + \dots$$

Furthermore, he can use exactly the same rule, Eq. 7.13, to compute the expectation value.

When the density matrix corresponds to a single state, it is a projection operator that projects onto that state. In this case, we say that the state is *pure*. A pure state represents the maximum amount of knowledge that Bob can have of a quantum system. But in the more general case, the density matrix is a mix of several projection operators. We then say that the density matrix represents a *mixed* state.

I have used the term *density matrix*, but strictly speaking, ρ is an operator. It only becomes a matrix when a basis is chosen. Suppose we choose the basis $|a\rangle$. The density matrix is just the matrix representation of ρ with respect to this basis:

$$\rho_{aa'} = \langle a | \rho | a' \rangle.$$

If the matrix representation of \mathbf{L} is $L_{a',a}$ then 7.13 takes the form

$$\langle \mathbf{L} \rangle = \sum_{a,a'} L_{a',a} \rho_{a,a'}. \quad (7.14)$$

7.4 Entanglement and Density Matrices

Classical physics also has its notion of pure and mixed states, although they are not called by those names. Just to illustrate, let's consider a system of two particles moving along a line. According to the rules of classical mechanics, we can calculate the orbits of the particles if we know the values of their positions (x_1 and x_2) and momenta (p_1 and p_2) at a certain instant in time. The state of the system is thus specified by four numbers: x_1 , x_2 , p_1 , and p_2 . If we know these four numbers, we have as complete a description of the two-particle system as it is possible to have: there is no more to know. We can call this a pure classical state.

Often, however, we don't know the exact state, but only some probabilistic information. That information can be encoded in a probability density

$$\rho(x_1, x_2, p_1, p_2).$$

A classical pure state is just a special case of a probability

density, in which ρ is nonzero at only one point. But more generally, ρ will be smeared out, in which case we could call it a classical mixed state.⁵ When ρ is smeared out, it means our knowledge of the system state is incomplete. The more smeared out it is, the greater our ignorance.

One thing should be completely obvious from this example: if you know the pure state for the combined two-particle system, then you know everything about each particle. In other words, a pure state for two *classical* particles implies a pure state for each of the individual particles.

But this is exactly what is *not* true in quantum mechanics when a system is entangled. The state of a composite system can be absolutely pure, but each of its constituents must be described by a mixed state.

Let's take a system composed of two parts, A and B . It could be two spins or any other composite system.

In this case, we will suppose that Alice has complete knowledge of the state of the combined system. In other words, she knows the wave function

$$\Psi(a, b).$$

There is nothing missing from her knowledge of the combined system. Nevertheless, Alice is not interested in B . Instead, she wishes to find out as much as she can about A without looking at B . She selects an observable \mathbf{L} that belongs to A , and that does nothing to B when it acts. The rule for

⁵By *smeared out*, we mean that $\rho(x_1, x_2, p_1, p_2)$ will be nonzero for a range of values of its arguments, not just one value. The greater this range, the more smeared out ρ becomes.

calculating the expectation value of \mathbf{L} is

$$\langle \mathbf{L} \rangle = \sum_{ab, a'b'} \Psi^*(a'b') L_{a'b', ab} \Psi(ab). \quad (7.15)$$

So far, this is entirely general. However, if the observable \mathbf{L} is associated only with A , then it acts trivially on the b -index and we can write the expectation value as

$$\langle \mathbf{L} \rangle = \sum_{a, b, a'} \Psi^*(a'b) L_{a', a} \Psi(ab). \quad (7.16)$$

Now, Alice can summarize all of her knowledge, at least for the purpose of studying A , in terms of a matrix ρ :

$$\rho_{aa'} = \sum_b \Psi^*(a'b) \Psi(ab). \quad (7.17)$$

Surprisingly, Eq. 7.16 has exactly the same form as Eq. 7.14 for expectation value of a mixed state. Indeed, only in the very special case of a product state will ρ have the form of a projection operator. In other words, despite the fact that the composite system is described by a perfectly pure state, the subsystem A must be described by a mixed state.

There's a subtle point about our notation for density matrices that's worth noticing: in Eq. 7.17, the right-hand index of ρ , that is, the a' index, corresponds to the complex conjugate state-vector $\Psi^*(a'b)$ in the summation. This is a consequence of our convention

$$L_{aa'} = \langle a | \mathbf{L} | a' \rangle$$

for labeling the matrix elements of an operator \mathbf{L} . Applying this convention to

$$\rho = |\Psi\rangle\langle\Psi|$$

results in

$$\rho_{aa'} = \langle a|\Psi\rangle\langle\Psi|a'\rangle,$$

or

$$\rho_{aa'} = \Psi(a)\Psi^*(a').$$

7.5 Entanglement for Two Spins

Before leading you further into the world of entanglement, I'll give you a simple definition and a quick warm-up exercise. If Alice only has a single spin in a known state, her density matrix is defined to be

$$\rho_{aa'} = \psi^*(a')\psi(a).$$

This equation tells you how to calculate an element of Alice's density matrix. If we stick with our familiar σ_z basis, each index a and a' can take the values *up* and *down*, so Alice has a 2×2 density matrix.

Exercise 7.4: Calculate the density matrix for

$$|\Psi\rangle = \alpha|u\rangle + \beta|d\rangle.$$

Answer:

$$\psi(u) = \alpha; \quad \psi^*(u) = \alpha^*$$

$$\psi(d) = \beta; \quad \psi^*(d) = \beta^*$$

$$\rho_{a'a} = \begin{pmatrix} \alpha^*\alpha & \alpha^*\beta \\ \beta^*\alpha & \beta^*\beta \end{pmatrix}.$$

Now try plugging in some numbers for α and β . Make sure they are normalized to 1. For example, $\alpha = \frac{1}{\sqrt{2}}$, $\beta = \frac{1}{\sqrt{2}}$.

This simple example is a good way to understand the properties of density matrices. You can refer back to it as we look at the more complex example of an entangled state.

Suppose we know the wave function of a composite system, for example

$$\psi(a, b),$$

but we are only interested in Alice's subsystem. In other words, we want to keep track of everything that Alice can ever measure. Do we have to know the whole wave function? Or is there some way to get rid of Bob's variables? The answer to the latter question is yes; we can capture Alice's complete description in terms of a density matrix ρ .

Let's consider an observable \mathbf{L} of Alice's system. Like any observable, it can of course be represented as a matrix:

$$L_{a'b',ab} = \langle a'b' | \mathbf{L} | ab \rangle.$$

Remember, for the composite system, the pair ab is really a single index labeling a basis vector.

When we say, “ \mathbf{L} is an Alice-observable,” what we mean is that \mathbf{L} does nothing to Bob’s half of the state-label. This forces some restrictions on the form of \mathbf{L} . The idea is to filter out (set equal to zero) any of \mathbf{L} ’s matrix elements that have the effect of changing Bob’s half of the state-label. In other words, \mathbf{L} has the special form

$$L_{a'b',ab} = L_{a'a} \delta_{b'b}. \quad (7.18)$$

This simple-looking equation requires some explanation, and you may want to review the material on tensor products in component form, in the Interlude on tensor products (Section 6.1). The left-hand side of the equation is an element of a 4×4 matrix. Each of its two indices can take four distinct values: uu , ud , du , or dd . What about the right-hand side? The matrix element $L_{a'a}$ also has two indices, but each of them can take only two distinct values: u or d . In fact, the same symbol L refers to two different matrices on each side of Eq. 7.18.

At first glance, it appears as though we have equated a 4×4 matrix to a 2×2 matrix, and indeed that would be a problem. However, the factor $\delta_{b'b}$ makes everything work out. The term $L_{a'a} \delta_{b'b}$ is an element of the tensor product of two 2×2 matrices, and that tensor product *is* a 4×4

matrix.⁶ Here is the way to read Eq. 7.18:

The 4×4 matrix $L_{a'b',ab}$ can be factored into a tensor product of the two 2×2 matrices $L_{a'a}$ and $\delta_{b'b}$, where $\delta_{b'b}$ is equivalent to the 2×2 identity matrix.

Now, let's calculate the expectation value of \mathbf{L} (the 4×4 version) using the full apparatus of the composite system:

$$\langle \Psi | \mathbf{L} | \Psi \rangle = \sum_{a,b,a',b'} \psi^*(a', b') L_{a'b',ab} \psi(a, b).$$

As I warned, there are lots of indices. But it gets simpler if we use the special form of the matrix L . The factor $\delta_{b'b}$ in Eq. 7.18—a Kronecker delta—filters out any elements that change Bob's half of the label, and leaves the others intact. It tells us to set $b' = b$ to get

$$\langle \Psi | \mathbf{L} | \Psi \rangle = \sum_{a',b,a} \psi^*(a', b) L_{a',a} \psi(a, b). \quad (7.19)$$

For the moment, let's ignore the sums over a and a' , and concentrate instead on the sum over b . We encounter the quantity

$$\rho_{a'a} = \sum_b \psi^*(a, b) \psi(a', b). \quad (7.20)$$

The 2×2 matrix $\rho_{a'a}$ is Alice's density matrix. Notice that $\rho_{a'a}$ does not depend on any b -index since it has already been

⁶We could also call it a Kronecker product, since we're talking about matrices. The formal distinction is not important for our purposes.

summed over b . It is purely a function of Alice variables a and a' . In fact, we only kept the b 's in the equation to make the example in the next section easier to follow.

We can simplify Eq. 7.19 by plugging in $\rho_{a'a}$ from Eq. 7.20. The expectation value of \mathbf{L} (the 2×2 version) then becomes

$$\langle \mathbf{L} \rangle = \sum_{a'a} \rho_{a'a} L_{a,a'}. \quad (7.21)$$

In summing over b , we have collapsed a 4×4 matrix down to a 2×2 matrix. This makes sense. We expect an operator that acts on the composite system to be a 4×4 matrix, and we expect an Alice operator to be 2×2 .

Notice that the right side of Eq. 7.21 is a sum of diagonal matrix elements. In other words, it's the trace of the matrix $\rho \mathbf{L}$, which we can write as

$$\langle \mathbf{L} \rangle = \text{Tr } \rho \mathbf{L}.$$

The lesson is this: To calculate Alice's density matrix ρ , we may need to know the full wave function, including the dependence on Bob's variables. But once we know ρ , we can forget where it came from, and use it to calculate *anything* about Alice's observations. As a simple example, we can use ρ to calculate the probability $P(a)$ that Alice's system will be left in the state a if a measurement is performed. To determine $P(a)$, we begin with $P(a, b)$, the probability that the combined system is in state $|ab\rangle$. That's just

$$P(a, b) = \psi^*(a, b)\psi(a, b).$$

By the standard rules of probability, if we sum over b , we get the probability for a :

$$P(a) = \sum_b \psi^*(a, b)\psi(a, b).$$

This is just a diagonal entry in the density matrix:

$$P(a) = \rho_{aa}. \quad (7.22)$$

Here are some properties of density matrices:

- Density matrices are Hermitian:

$$\rho_{aa'} = \rho_{a'a}^*.$$

- The trace of a density matrix is 1:

$$\text{Tr}(\rho) = 1.$$

Eq. 7.22 should help make this clear because the left side is a probability.

- The eigenvalues of the density matrix are all positive and lie between 0 and 1. It follows that if any eigenvalue is 1, all the others are 0. Can you interpret this result?
- For a pure state:

$$\begin{aligned} \rho^2 &= \rho \\ \text{Tr}(\rho^2) &= 1 \end{aligned}$$

- For a mixed or entangled state:

$$\rho^2 \neq \rho$$

$$\text{Tr}(\rho^2) < 1$$

The last two properties give us a clear way to distinguish mathematically between pure and mixed states. A subsystem of an entangled state (such as Alice’s half of the singlet state) is considered a mixed state.

It’s worth taking a moment to understand these two properties a little better. To simplify things, we will assume that ρ is a diagonal matrix—in other words, all of its off-diagonal elements are zero. This simplification costs us nothing because ρ is Hermitian, and it turns out that every Hermitian matrix can be expressed in diagonal form in some basis.⁷ Taking the square of a diagonal matrix is quite simple: all you need to do is square each individual element. Since ρ represents a mixed state, and the diagonal elements of ρ must add up to 1, *none* of the diagonal elements of ρ can *equal* 1. Otherwise, ρ would represent a pure state. Therefore, ρ must have at least two positive diagonal elements that are less than 1. Squaring these elements gives a new matrix ρ^2 whose elements are even smaller. This accounts for both of the mixed-state properties of ρ .

Before you try the next exercises, I’ll mention one more thing about the trace. It turns out that the trace has many

⁷As we mentioned earlier, in Section 7.2, a Hermitian matrix \mathbf{M} can be diagonalized by a transformation $\mathbf{P}^\dagger \mathbf{M} \mathbf{P}$, where \mathbf{P} is a unitary matrix whose columns are the normalized eigenvectors of \mathbf{M} .

interesting mathematical properties. One of its more useful properties is that the trace of a product of two matrices does not depend on their order of multiplication. In other words,

$$\text{Tr}AB = \text{Tr}BA,$$

even if

$$AB \neq BA.$$

I mention this because you will sometimes see the trace of the density matrix written as $\text{Tr} \mathbf{L}\rho$, instead of $\text{Tr} \rho\mathbf{L}$. These two expressions are equivalent.

Exercise 7.5:

a) Show that

$$\begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix}^2 = \begin{pmatrix} a^2 & 0 \\ 0 & b^2 \end{pmatrix}.$$

b) Now, suppose

$$\rho = \begin{pmatrix} \frac{1}{3} & 0 \\ 0 & \frac{2}{3} \end{pmatrix}.$$

Calculate

$$\begin{aligned} &\rho^2 \\ &\text{Tr}(\rho) \\ &\text{Tr}(\rho^2). \end{aligned}$$

c) If ρ is a density matrix, does it represent a pure state or a mixed state?

Exercise 7.6: Use Eq. 7.22 to show that if ρ is a density matrix, then

$$\text{Tr}(\rho) = 1.$$

7.6 A Concrete Example: Calculating Alice's Density Matrix

So far, the discussion of density matrices may have been a little abstract for some readers. Here is a worked-out example that should help bring density matrices into sharper focus. Recall the definition of Alice's density matrix from Eq. 7.20:

$$\rho_{a'a} = \sum_b \psi^*(a, b) \psi(a', b). \quad (7.23)$$

Now, consider the state-vector

$$|\Psi\rangle = \frac{1}{\sqrt{2}}(|ud\rangle + |du\rangle).$$

Notice that two of the basis vectors have a coefficient of $\frac{1}{\sqrt{2}}$, while the other two have coefficients of zero. The state is normalized because the sum of the squared coefficients is 1. Also, all four coefficients happen to be real, which simplifies the process of complex conjugation.

Let's calculate Alice's density matrix for this state. First, for all possible inputs a and b , we'll list the values of $\psi(a, b)$. Recall that these are just the basis vector coefficients:

$$\psi(u, u) = 0$$

$$\psi(u, d) = \frac{1}{\sqrt{2}}$$

$$\psi(d, u) = \frac{1}{\sqrt{2}}$$

$$\psi(d, d) = 0.$$

Next, we'll use these four equations to calculate each element of Alice's density matrix by expanding the summation of Eq. 7.23. In the expansion, notice that for every factor of the form $\psi^*(a, b)\psi(a', b)$, Bob's input is the same for both factors. We discard any terms that do not have this property. This is what we mean by "setting b' equal to b in the summation." Here is the expansion:

$$\rho_{uu} = \psi^*(u, u)\psi(u, u) + \psi^*(u, d)\psi(u, d) = \frac{1}{2}$$

$$\rho_{ud} = \psi^*(u, u)\psi(d, u) + \psi^*(u, d)\psi(d, d) = 0$$

$$\rho_{du} = \psi^*(d, u)\psi(u, u) + \psi^*(d, d)\psi(u, d) = 0$$

$$\rho_{dd} = \psi^*(d, u)\psi(d, u) + \psi^*(d, d)\psi(d, d) = \frac{1}{2}.$$

These values are the elements of a 2×2 matrix:

$$\rho = \begin{pmatrix} \frac{1}{2} & 0 \\ 0 & \frac{1}{2} \end{pmatrix}. \quad (7.24)$$

The trace of our matrix is 1. And our density matrix is done.⁸

Exercise 7.7: Use Eq. 7.24 to calculate ρ^2 . How does this result confirm that ρ represents an entangled state? We'll soon discover that there are other ways to check for entanglement.

Exercise 7.8: Consider the following states:

$$\begin{aligned} |\psi_1\rangle &= \frac{1}{2}(|uu\rangle + |ud\rangle + |du\rangle + |dd\rangle) \\ |\psi_2\rangle &= \frac{1}{\sqrt{2}}(|uu\rangle + |dd\rangle) \\ |\psi_3\rangle &= \frac{1}{5}(3|uu\rangle + 4|ud\rangle). \end{aligned}$$

For each one, calculate Alice's density matrix and Bob's density matrix. Check their properties.

7.7 Tests for Entanglement

Suppose I gave you a wave function

$$\psi(a, b)$$

⁸Art's a poet, and he's not even aware of it.

for the composite S_{AB} system. How could you tell whether the corresponding state is entangled? I am not referring to an experimental test but to a mathematical procedure. A related question is whether there are varying degrees of entanglement. If there are, how could you quantify them?

Entanglement is the quantum mechanical generalization of correlation. In other words, it indicates that Alice can learn something about Bob's half of the system by measuring her own. In the classical example of the previous lecture, I illustrated the idea of correlation using coins. If Alice observes the coin that Charlie gave her, she not only knows whether her own coin is a penny or a dime; she also knows which coin Bob has. That's the experimental picture. The *mathematical* indication of correlation is that the probability function $P(a, b)$ does not factorize (that is, it does not look like Eq. 6.3). Whenever the probability distribution does not factorize, there are nonzero correlations as I described in Inequality 6.2.

7.7.1 The Correlation Test for Entanglement

Let's assume that \mathbf{A} is an Alice observable and \mathbf{B} is a Bob observable. The correlation between them is defined in terms of the average values (also known as the expectation values) of the individual observables, and of their product. Suppose that

$$\langle \mathbf{A} \rangle$$

$$\langle \mathbf{B} \rangle$$

$$\langle \mathbf{AB} \rangle$$

are these expectation values. The correlation $C(A, B)$ between \mathbf{A} and \mathbf{B} is defined as

$$C(A, B) = \langle \mathbf{AB} \rangle - \langle \mathbf{A} \rangle \langle \mathbf{B} \rangle.$$

Exercise 7.9: Given any Alice observable \mathbf{A} and Bob observable \mathbf{B} , show that for a product state, the correlation $C(A, B)$ is zero.

From this exercise, we can learn something about entanglement. If a system is in a state where one can find any two observables \mathbf{A} and \mathbf{B} that are correlated—meaning that $C(A, B) \neq 0$ —then the state is entangled. Correlations are defined to lie in the range -1 to $+1$. These extreme values represent the greatest possible negative and positive correlations. The greater the magnitude of $C(A, B)$, the more entangled is the state. If $C(A, B) = 0$, then there is no correlation (and no entanglement) at all.

7.7.2 The Density Matrix Test for Entanglement

To calculate correlations, you have to know about both Bob's part and Alice's part of the system, along with the system wave function. But there is another test for entanglement

that only requires us to know Alice's (or Bob's) density matrix. Let's suppose that the state $|\Psi\rangle$ is a product state of a Bob factor $|\phi\rangle$ and an Alice factor $|\psi\rangle$. That means the composite wave function is also the product of a Bob factor and an Alice factor:

$$\psi(a, b) = \psi(a)\phi(b).$$

Now, let's work out Alice's density matrix. We use the definition in Eq. 7.20 to get

$$\rho_{a'a} = \psi^*(a)\psi(a') \sum_b \phi^*(b)\phi(b).$$

But if Bob's state is normalized, then

$$\sum_b \phi^*(b)\phi(b) = 1,$$

which makes Alice's density matrix particularly simple:

$$\rho_{a'a} = \psi^*(a)\psi(a'). \tag{7.25}$$

Notice that it only depends on the Alice variables. Perhaps it's not very surprising that everything we need to know about Alice's system is contained in Alice's wave function.

Now, I'm going to prove a key theorem about the eigenvalues of Alice's density matrix, under the assumption of a product state. It is true only for unentangled states and serves to identify them. The theorem says that for any product state, Alice's (or Bob's) density matrix has exactly one

nonzero eigenvalue, and that eigenvalue is exactly 1. We begin the theorem by writing the eigenvalue equation for the matrix ρ :

$$\sum_{a'} \rho_{a'a} \alpha_{a'} = \lambda \alpha_a.$$

In other words, the matrix ρ acting on the column vector α gives back the same vector multiplied by an eigenvalue λ . Using the simple form of ρ in Eq. 7.25, we can write

$$\psi(a') \sum_a \psi^*(a) \alpha_a = \lambda \alpha_{a'}. \quad (7.26)$$

Now, you may notice a couple of things. First, the quantity

$$\sum_a \psi^*(a) \alpha_a$$

has the form of an inner product. If the column vector α is orthogonal to ψ , then the left side of Eq. 7.26 is zero. Such a vector is an eigenvector of ρ with eigenvalue zero.

If the dimension of Alice's space of states is N_A , then there are $N_A - 1$ vectors orthogonal to ψ . Each one of them is an eigenvector of ρ with eigenvalue 0. That leaves only one possible direction for an eigenvector with a nonzero eigenvalue, namely the vector $\psi(a)$. In fact, if we plug in $\alpha_a = \psi(a)$, we do indeed find that it is an eigenvector of ρ with eigenvalue 1.

To summarize the theorem: If the composite Alice-Bob system is in a product state, then Alice's (or Bob's) density

matrix has one and only one eigenvalue equal to 1, and all the rest are zero. Moreover, the eigenvector with a nonzero eigenvalue is nothing but the wave function of Alice's half of the system.

In this situation, Alice's system is in a pure state. All of Alice's observations are described as if Bob and his system never existed and Alice had an isolated system described by the wave function $\psi(a')$.

The opposite extreme of a pure state is a *maximally entangled* state. Maximally entangled states are states of a combined system in which nothing is known about either subsystem, even though they are complete descriptions of the system as a whole—as complete as quantum mechanics allows. The state $|sing\rangle$ is a maximally entangled state.

When Alice calculates her density matrix for a maximally entangled state, she finds something very disappointing: the density matrix is proportional to the unit matrix. All the eigenvalues are equal, and given that they all sum to unity, each eigenvalue is equal to $1/N_A$. In other words,

$$\rho_{a'a} = \frac{1}{N_A} \delta_{a'a}. \quad (7.27)$$

Why is Alice disappointed? Go back to Eq. 7.22. This equation says that the probability for a particular state a is the diagonal element of ρ , but Eq. 7.27 tells us that all the probabilities are equal. What could be less informative than a probability distribution so structureless that every possible outcome is equally probable?

Maximal entanglement implies a complete lack of information about Alice's subsystem for experiments that only involve that one subsystem. On the other hand, it implies a large correlation between Alice's and Bob's measurements. For the singlet state, if Alice measures any component of her spin, she automatically knows the result Bob would get if he were to measure the same component of his spin. This is exactly the kind of knowledge that is precluded in a product state.

So in each type of state, some things are predictable and some are not. In a product state, we can make statistical predictions about measurements made on each separate subsystem, but Alice's measurements tell her nothing about Bob's system. In a maximally entangled state, on the other hand, Alice can predict nothing about her own measurements, but she knows a great deal about the relation between her outcomes and Bob's.

7.8 The Process of Measurement

We have seen that quantum systems evolve in what look like irreconcilably different ways: by unitary evolution between measurements, and by wave function collapse when measurements take place. This circumstance has led to some of the most contentious debates and confusing claims about so-called reality. I'm going to steer away from those debates and stick to the facts. Once you know how quantum mechanics works, you can decide for yourself whether you think there is a problem.

Let's begin by noting that every measurement involves

a system and an apparatus. But if quantum mechanics is a consistent theory, then it should be possible to combine the system and apparatus into a single bigger system. For simplicity let's take the system to be a single spin. The apparatus \mathcal{A} is the same one that we used in the very first lecture. The window in the apparatus can show three possible readings. The first is blank—it represents the neutral state of the apparatus before it comes in contact with the spin. The two other readings record the two possible outcomes of the measurement: $+1$ or -1 .

If the apparatus is a quantum system (of course, it must be), then it is described by a space of states. In the simplest description, the apparatus has exactly three states: a blank state and two outcome states. Thus, the basis vectors for the apparatus are

$$\begin{aligned} &|b\rangle \\ &|+1\rangle \\ &|-1\rangle. \end{aligned}$$

Meanwhile, the basis states of the spin can be taken to be the usual *up* and *down* states:

$$\begin{aligned} &|u\rangle \\ &|d\rangle. \end{aligned}$$

From these two sets of basis vectors, we can build up a composite (tensor product) space of states that has the six basis vectors

$$\begin{aligned}
&|u, b\rangle \\
&|u, +1\rangle \\
&|u, -1\rangle \\
&|d, b\rangle \\
&|d, +1\rangle \\
&|d, -1\rangle.
\end{aligned}$$

The detailed mechanics of what takes place when system meets apparatus may be complicated, but we are free to make some assumptions about how the combined system evolves. Let's assume the apparatus starts in the blank state and the spin starts in the *up* state. After the apparatus interacts with the spin, the final state (by assumption) is

$$|u, +1\rangle.$$

In other words, the interaction leaves the spin unchanged but flips the apparatus to the $+1$ state. We write this as

$$|u, b\rangle \rightarrow |u, +1\rangle. \quad (7.28)$$

Similarly, we can require that if the spin is in the *down* state, it flips the apparatus to the -1 state:

$$|d, b\rangle \rightarrow |d, -1\rangle. \quad (7.29)$$

So by looking at the apparatus after it interacts with the spin, you can tell what the spin was initially. Now, let's assume that the initial spin state is more general, namely

$$\alpha_u|u\rangle + \alpha_d|d\rangle.$$

If we include the apparatus as part of the system, the initial state is

$$\alpha_u|u, b\rangle + \alpha_d|d, b\rangle. \quad (7.30)$$

This initial state is a product state, specifically a product of the initial spin state and the blank apparatus state. You can check that it is completely unentangled.

Exercise 7.10: Verify that the state-vector in 7.30 represents a completely unentangled state.

Because we know from Eqs. 7.28 and 7.29 how the individual terms in 7.30 evolve, we can easily determine the final state:

$$\alpha_u|u, b\rangle + \alpha_d|d, b\rangle \rightarrow \alpha_u|u, +1\rangle + \alpha_d|d, -1\rangle.$$

This final state is an entangled state. In fact, if $\alpha_u = -\alpha_d$, it is the maximally entangled singlet state. Indeed, one can look at the apparatus and immediately tell what the spin state is: if the apparatus reads $+1$, the spin is *up*, and if it reads -1 , the spin is *down*. Moreover, the probability that the final apparatus shows $+1$ is

$$\alpha_u^* \alpha_u.$$

This number represents a probability—it's exactly the same as the original probability that the spin was up. In this description of a measurement, no collapse of the wave function takes place. Instead, entanglement between the apparatus and the system just happens by unitary evolution of the state-vector.

The only problem is that, in a certain sense, we have merely delayed the difficulty. It is not very satisfying to be told that the apparatus “knows” the spin state unless the experimenter—let's say Alice—is allowed to look at the apparatus. Isn't it true that when she does so, she will collapse the wave function of the composite system? Yes and no. For all of Alice's purposes, yes; she will conclude that the apparatus, and the spin, are in one of the two possible configurations and will proceed accordingly.

But now let's bring Bob into the picture. So far, he has not interacted with the spin, the apparatus, or Alice. From his point of view, all three form a single quantum system. No wave function collapse took place when Alice looked at the apparatus. Instead, Bob says that Alice became entangled with the other two component systems.

That's all well and good, but what happens when Bob looks at Alice? For his purposes, he has collapsed the wave function. But then there is good old Charlie . . .

Does the last entity to look at the system collapse the wave function, or does it just get entangled? Or is there

a last looker? I won't try to answer these questions, but what should be apparent is that quantum mechanics is a consistent calculus of probabilities for a certain kind of experiment involving a system and an apparatus. We use it, and it works, but when we try to ask questions about the underlying "reality," we get confused.

7.9 Entanglement and Locality

Does quantum mechanics violate locality? Some people think so. Einstein railed against the "spooky action at a distance" (*spukhafte Fernwirkung*) that he claimed was implied by quantum mechanics. And John Bell became almost a cult figure by proving that quantum mechanics is nonlocal.

On the other hand, most theoretical physicists, particularly those who study quantum field theory, which is riddled with entanglement, would claim the opposite: quantum mechanics done correctly ensures locality.

The problem, of course, is that the two groups mean different things by *locality*. Let's begin with the quantum field theorist's understanding of the term. From this point of view, locality has only one meaning: it is impossible to send a signal faster than the speed of light. I will show you how quantum mechanics enforces this rule.

First, let me expand the definition of *Alice's system* and *Bob's system*. So far, I have used the term *Alice's system* to mean some system that Alice carries with her and can do experiments on. For the rest of this section, I will use the term to mean something else: Alice's system consists not only of some system that she carries, but also the apparatus

that she uses, and even herself. The same thing, of course, goes for Bob's system. The basis ket-vectors

$$|a\rangle$$

describe everything that Alice can interact with. Likewise, the ket-vectors

$$|b\rangle$$

describe everything that Bob can interact with. And the tensor product states

$$|ab\rangle$$

describe the combination of Alice's and Bob's worlds.

We will assume that Alice and Bob may have been close enough to interact sometime in the past, but at present Alice is on Alpha Centauri and Bob is in Palo Alto. The Alice-Bob wave function is

$$\psi(ab),$$

and it may be entangled. Alice's complete description of her system, her apparatus, and herself is contained in her density matrix ρ :

$$\rho_{aa'} = \sum_b \psi^*(a'b) \psi(ab). \quad (7.31)$$

Consider this question: Can Bob, at his end, do anything to instantly change Alice’s density matrix? Keep in mind that Bob can only do things that the laws of quantum mechanics allow. In particular, Bob’s evolution, whatever causes it, must be unitary. In other words, it must be described by a unitary matrix

$$\mathbf{U}_{bb'}.$$

The matrix \mathbf{U} represents whatever happens to Bob’s system, whether or not Bob does an experiment. It acts on the wave function to produce a new wave function, which we’ll call the “final” wave function:

$$\psi_{final}(ab) = \sum_{b'} \mathbf{U}_{bb'} \psi(ab').$$

We can also write the complex conjugate of this wave function:

$$\psi_{final}^*(a'b) = \sum_{b''} \psi^*(a'b'') \mathbf{U}_{b''b}^\dagger.$$

Notice that we added primes to some of the symbols to avoid mixing them up in the next step. Now, let’s calculate Alice’s new density matrix. We’ll use Eq. 7.31, but we’ll replace the original wave functions with the final ones:

$$\rho_{aa'} = \sum_{b,b',b''} \psi^*(a'b'') \mathbf{U}_{b''b}^\dagger \mathbf{U}_{bb'} \psi(ab').$$

There are lots of indices flying around now, but the math isn't as hard as it looks. In fact, look at how the \mathbf{U} matrices enter through the combination

$$\mathbf{U}_{b''b}^\dagger \mathbf{U}_{bb'}.$$

This combination is just the matrix product $\mathbf{U}^\dagger \mathbf{U}$. But recall that \mathbf{U} is unitary. This tells you that the product $\mathbf{U}^\dagger \mathbf{U}$ is the unit matrix $\delta_{b''b'}$. As before, this amounts to an instruction to include all the terms where $b'' = b'$, and to ignore all the others. With this simplification, we get

$$\rho_{aa'} = \sum_b \psi^*(a'b) \psi(ab).$$

This is exactly the same as Eq. 7.31. In other words, $\rho_{aa'}$ is exactly the same as it was before \mathbf{U} acted. Nothing that happens at Bob's end has any immediate effect on Alice's density matrix, even if Bob and Alice are maximally entangled. This means that Alice's view of her subsystem (her statistical model) remains exactly as it was. This remarkable result may seem surprising for a maximally entangled system, but it also guarantees that no faster-than-light signal has been sent.

7.10 The Quantum Sim: An Introduction to Bell's Theorem

It's interesting that unitarity played a prominent role in guaranteeing that no signal can be sent instantaneously. If \mathbf{U} had not been unitary, Alice's final density matrix *would indeed* have been affected by Bob.

What was it, then, that disturbed Einstein so much that he spoke of spooky action at a distance? To answer this question, it's important to understand that he and Bell were talking about a totally different notion of locality. To illustrate this, I am going to invent a computer game. What my new computer game does is try to fool you into thinking there is a quantum spin in a magnetic field inside the computer. You get to do experiments to test this possibility. See Fig. 7.1 for a schematic.

Here's how it works: Inside the computer, the memory stores two complex numbers, α_u and α_d , subject to the usual normalization rule,

$$\alpha_u^* \alpha_u + \alpha_d^* \alpha_d = 1.$$

At the beginning of the game, the α coefficients are initialized at some value. The computer then solves the Schrödinger equation to update the α 's exactly as if they were the components of the spin's state-vector.

The computer also stores the classical three-dimensional orientation of the apparatus in the form of two angles or a

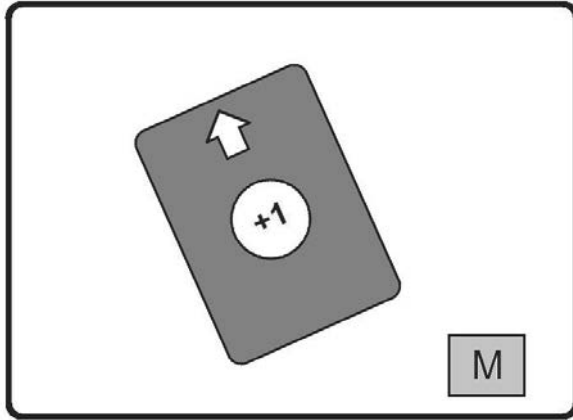


Figure 7.1: Quantum Sim. The computer screen displays the user-controlled orientation of the apparatus. For simplicity, only the two-dimensional orientation is shown here. The user can press the M button whenever she or he wants to measure the spin (not shown). Between measurements, the spin state evolves according to the Schrödinger equation.

unit vector. The keyboard allows you to set these angles and change them at will. One more element is stored in the memory, namely the value (either $+1$ or -1) representing the number in the window of the apparatus. The computer screen shows the apparatus. As the experimenter, you get to choose how your apparatus will be oriented. There is also

a measure button M that activates the apparatus.

The final element of the program is a random number generator that produces the measurement results $+1$ or -1 with probabilities $\alpha_u^* \alpha_u$ and $\alpha_d^* \alpha_d$, respectively. Keep in mind that random number generators are not really generators of random numbers; they are random number simulators. They are based on entirely classical deterministic mechanisms, using things like the digits of π to generate numbers. Nevertheless, they are good enough to fool you.

The game begins, and the computer continually updates the values of α_u and α_d . You wait as long as you want and then hit the M button. Then, with the aid of the random number generator, the game produces an outcome that is displayed on the screen. Based on this outcome, the computer updates the state by collapse. If the outcome is $+1$, the value of α_d is reset to zero, and the value of α_u is reset to unity. If the outcome is -1 , the value of α_d is reset to unity, and the value of α_u is reset to zero. Then, the Schrödinger equation takes over until you hit M again.

Being a good experimenter, you do many trials and collect statistics, which you compare with quantum mechanical predictions. If everything works properly, you conclude that quantum mechanics is the correct description of whatever is taking place in the computer. Of course, the computer is still entirely classical, but it simulates a quantum spin without much difficulty.

Next, let's try the same thing with two computers, A and B, simulating two quantum spins. If the spins start in a product state and never interact, we can simply play

the game on each of the two computers without any cross talk. But now, Alice, Bob, and Charlie return to help us out. Charlie, of course, wants to create an entangled pair. He begins by connecting the two computers with a cable to form a single computer, and we assume the cable can send instantaneous signals. In its memory, the combined computer now stores four complex numbers,

$$\alpha_{uu}, \alpha_{ud}, \alpha_{du}, \alpha_{dd},$$

and it updates these numbers using the Schrödinger equation. Each computer screen shows an apparatus. Alice's screen shows \mathcal{A} and Bob's screen shows \mathcal{B} . Each virtual apparatus can be independently oriented, and each can be independently activated by its own M button. When either M button is pressed, the joint memory (with the aid of the random number generator) sends a signal to the corresponding apparatus and produces an outcome.

Can this device simulate the quantum mechanics of the two-spin system? Yes, it can—as long as the cable connecting the computers is not disconnected, and as long as it can send messages instantaneously. But unless the system is in a product state and stays in a product state, disconnecting the two computers will destroy the simulation.

Can we prove this? Again, the answer is yes—and that is the essential content of Bell's theorem. Any classical simulation of quantum mechanics that tries to spatially separate Alice's and Bob's apparatuses must have an instantaneous cable connecting the separate computers with a central memory that stores and updates the state-vector.

But doesn't this mean that locality-violating information can be sent through the cable? It would, if Alice, Bob, and Charlie were allowed to do anything that nonrelativistic classical systems can do.⁹ But if the only operations that are allowed are those that simulate quantum operations, then the answer is no. As we've seen, quantum mechanics does not allow Alice's density matrix to be affected by Bob's actions.

This problem is not a problem for quantum mechanics. It's a problem for *simulating* quantum mechanics with a classical Boolean computer. That's the content of Bell's theorem: The classical computers have to be connected with an instantaneous cable to simulate entanglement.

7.11 Entanglement Summary

Of all the counterintuitive ideas quantum mechanics forces upon us, entanglement may be the hardest one to accept. There is no classical analog for a system whose full state description contains no information about its individual sub-components. Nonlocality is surprisingly difficult to even define. The best way to come to terms with these issues is to internalize the mathematics. What follows is a compact summary of what we've learned about entanglement. In particular, we've tried to map out the differences between entangled, unentangled, and partially entangled states by creating "rap sheets" for three specific examples—the singlet state, a product state, and a "near singlet" state. We hope this

⁹In other words, systems that permit signals to be sent instantly.

format will help clarify the mathematical similarities and differences. Please take some time to review this material and work the exercises before moving on.

State-Vector Rap Sheet 1

Name: Product State (No Entanglement)

Wanted for: Excessive Locality, Impersonating a Classical System

Description: Each subsystem is fully characterized. There are no correlations between Alice's and Bob's systems.

State-Vector: $\alpha_u\beta_u|uu\rangle + \alpha_u\beta_d|ud\rangle + \alpha_d\beta_u|du\rangle + \alpha_d\beta_d|dd\rangle$

Normalization: $\alpha_u^*\alpha_u + \alpha_d^*\alpha_d = 1$, $\beta_u^*\beta_u + \beta_d^*\beta_d = 1$

Density Matrix: Alice's density matrix has exactly one nonzero eigenvalue, which equals 1. The eigenvector with this nonzero eigenvalue is the wave function of Alice's subsystem. The same goes for Bob.

Wave Function: Factorized: $\psi(a)\phi(b)$

Expectation Values:

$$\langle\sigma_x\rangle^2 + \langle\sigma_y\rangle^2 + \langle\sigma_z\rangle^2 = 1$$

$$\langle\tau_x\rangle^2 + \langle\tau_y\rangle^2 + \langle\tau_z\rangle^2 = 1$$

Correlation: $\langle\sigma_z\tau_z\rangle - \langle\sigma_z\rangle\langle\tau_z\rangle = 0$

State-Vector Rap Sheet 2**Name:** Singlet State (Maximum Entanglement)**Wanted for:** Nonlocality, Complete Quantum Weirdness**Description:** The composite system as a whole is fully characterized. There is no information about Alice's or Bob's subsystems.**State-Vector:** $\frac{1}{\sqrt{2}}(|ud\rangle - |du\rangle)$ **Normalization:** $\psi_{uu}^* \psi_{uu} + \psi_{ud}^* \psi_{ud} + \psi_{du}^* \psi_{du} + \psi_{dd}^* \psi_{dd} = 1$ **Density Matrix:**Full Composite System: $\rho^2 = \rho$, and $Tr(\rho^2) = 1$.Alice's Subsystem: Density matrix is proportional to the unit matrix, having equal eigenvalues that add up to 1. Hence, each measurement outcome is equally likely. $\rho^2 \neq \rho$, and $Tr(\rho^2) < 1$.**Wave Function:** Not Factorized: $\psi(a, b)$ **Expectation Values:**

$$\langle \sigma_z \rangle, \langle \sigma_x \rangle, \langle \sigma_y \rangle = 0$$

$$\langle \tau_z \rangle, \langle \tau_x \rangle, \langle \tau_y \rangle = 0$$

$$\langle \tau_z \sigma_z \rangle, \langle \tau_x \sigma_x \rangle, \langle \tau_y \sigma_y \rangle = -1$$

Correlation: $\langle \sigma_z \tau_z \rangle - \langle \sigma_z \rangle \langle \tau_z \rangle = -1$

State-Vector Rap Sheet 3

Name: “Near-Singlet” (Partial Entanglement)

Wanted for: Indecision, General Wishy-Washiness, Trouble Telling *up* from *down*

Description: There is some information about the composite system, and some about each subsystem. Incomplete in each case.

State-Vector: $\sqrt{0.6}|ud\rangle - \sqrt{0.4}|du\rangle$

Normalization: $\psi_{uu}^*\psi_{uu} + \psi_{ud}^*\psi_{ud} + \psi_{du}^*\psi_{du} + \psi_{dd}^*\psi_{dd} = 1$

Density Matrix:

Full Composite System: $\rho^2 \neq \rho$, and $Tr(\rho^2) < 1$.

Alice’s Subsystem: $\rho^2 \neq \rho$, and $Tr(\rho^2) < 1$.

Wave Function: Not Factorized: $\psi(a, b)$

Expectation Values:

$$\langle \sigma_z \rangle = 0.2$$

$$\langle \sigma_x \rangle, \langle \sigma_y \rangle = 0; \langle \tau_z \rangle = -0.2$$

$$\langle \tau_x \rangle, \langle \tau_y \rangle = 0$$

$$\langle \tau_z \sigma_z \rangle = -1$$

$$\langle \tau_x \sigma_x \rangle = -2\sqrt{0.24}$$

Correlation: $\langle \sigma_z \tau_z \rangle - \langle \sigma_z \rangle \langle \tau_z \rangle = -0.96$ for this example.

For partially entangled states in general, correlation is between -1 and $+1$, but not exactly 0 .

Exercise 7.11: Calculate Alice’s density matrix for σ_z for the “near-singlet” state.

Exercise 7.12: Verify the numerical values in each rap sheet.

Lecture 8

Particles and Waves

Art and Lenny have had enough entanglement for now. They're ready for something simpler.

Lenny: Hey Hilbert, do you have anything in one dimension?

Hilbert: Let me check. Single dimensions are very popular lately. Sometimes we run out.

Art: I'd settle for something classical, if that's all you have.

Hilbert: Not here, friend. We'd lose our license.

Art: Good point.

To the person in the street, quantum mechanics is all about light being particles and electrons being waves. But up until now, I've hardly mentioned particles, and the only mention of waves has been the wave function, which so far has had nothing to do with waves. So when do we get to the “real” quantum mechanics?

The answer, of course, is that real quantum mechanics is not so much about particles and waves as it is about the nonclassical logical principles that govern their behavior. Particle-wave duality is an easy extension of the things you've already learned, as we'll see in this lecture. But before we get into the physics, I want to review some mathematics, some of which is old—it appeared in earlier lectures—and some of which is new.

8.1 Mathematical Interlude: Working with Continuous Functions

8.1.1 Wave Function Review

We'll be using the language of wave functions in this lecture, so let's review some of that material before we dive in. In Lecture 5, we discussed wave functions as abstract objects, without explaining what they had to do with either waves or functions. Before correcting this omission, I will review what we discussed earlier.

Begin by picking an observable \mathbf{L} , with eigenvalues λ and eigenvectors $|\lambda\rangle$. Let $|\Psi\rangle$ be a state-vector. Since the eigenvectors of a Hermitian operator form a complete orthonormal basis, the vector $|\Psi\rangle$ can be expanded as

$$|\Psi\rangle = \sum_{\lambda} \psi(\lambda) |\lambda\rangle. \quad (8.1)$$

As you recall from Sections 5.1.2 and 5.1.3, the quantities

$$\psi(\lambda)$$

are called the wave function of the system. But notice: the specific form of $\psi(\lambda)$ depends on the specific observable \mathbf{L} that we initially choose. If we pick a different observable, the wave function (along with the basis vectors and eigenvalues) will be different, *even though we're still talking about the same state*. Therefore, we should qualify the statement that $\psi(\lambda)$ is the wave function associated with $|\Psi\rangle$. To be more precise, we should say that $\psi(\lambda)$ is the wave function in the \mathbf{L} -basis. If we use the orthonormality properties of the basis vectors,

$$\langle \lambda_i | \lambda_j \rangle = \delta_{ij},$$

then the wave function in the \mathbf{L} -basis may also be identified with the inner products (or projections) of the state-vector $|\Psi\rangle$ onto the eigenvectors $|\lambda\rangle$:

$$\psi(\lambda) = \langle \lambda | \Psi \rangle.$$

You can think of the wave function in two ways. First of all, it is the set of components of the state-vector in a particular basis. These components can be stacked up to form a column vector:

$$\begin{pmatrix} \psi(\lambda_1) \\ \psi(\lambda_2) \\ \psi(\lambda_3) \\ \psi(\lambda_4) \\ \psi(\lambda_5) \end{pmatrix}.$$

Another way to think of the wave function is as a function of λ . If you specify any allowable value of λ , the function $\psi(\lambda)$ produces a complex number. One can therefore say that

$$\psi(\lambda)$$

is a *complex-valued function of the discrete variable* λ . When thought of in this way, linear operators become operations that are applied to functions, and give back new functions.

One last reminder: the probability for an experiment to have outcome λ is

$$P(\lambda) = \psi^*(\lambda)\psi(\lambda).$$

8.1.2 Functions as Vectors

Up until now, the systems we have studied have had finite dimensional state-vectors. For example, the simple spin is described by a two-dimensional space of states. For this reason, the observables have had only a finite number of possible observable values. But there are more complicated observables that can have an infinite number of values. An example is a particle. The coordinates of a particle are observables, but, unlike spin, the coordinates have an infinite number of possible values. For instance, a particle moving along the x axis can be found at any real value of x . In other words, x is a continuously infinite variable. When the observables of a system are continuous, the wave function truly becomes a function of a continuous variable. To apply quantum mechanics to this kind of system, we have to expand the idea

of vectors to include functions.

Functions are functions, and vectors are vectors—they seem like different things, so in what sense are functions vectors? If you think of vectors as arrows pointing in three-dimensional space, then they are not the same as functions. But if you take the broader view of vectors as a set of mathematical objects satisfying certain postulates, then functions can indeed form a vector space. Such a vector space is often called a Hilbert space after the mathematician David Hilbert.

Let's consider the set of complex functions $\psi(x)$ of a single real variable x . By complex functions, I mean that for each x , $\psi(x)$ is a complex number. On the other hand, the independent variable x is an ordinary real variable. It can take on any real value from $-\infty$ to $+\infty$.

Now, let's nail down what we mean when we say "Functions are vectors." This is not a loose analogy or a metaphor. With appropriate restrictions (that we'll come back to), functions like $\psi(x)$ *satisfy the mathematical axioms that define a vector space*. We mentioned this idea briefly in Section 1.9.2, and now we'll make full use of it. Looking back at the axioms that define a complex vector space (in Section 1.9.1), we can see that complex functions satisfy all of them:

1. The sum of any two functions is a function.
2. The addition of functions is commutative.
3. The addition of functions is associative.
4. There is a unique *zero function* such that when you add it to any function, you get the same function back.

5. Given any function $\psi(x)$, there is a unique function $-\psi(x)$ such that $\psi(x) + (-\psi(x)) = 0$.
6. Multiplying a function by any complex number gives a function and is linear.
7. The distributive property holds, which means that

$$\begin{aligned}z[\psi(x) + \phi(x)] &= z\psi(x) + z\phi(x) \\[z + w]\psi(x) &= z\psi(x) + w\psi(x),\end{aligned}$$

where z and w are complex numbers.

All of this implies that we can identify the functions $\psi(x)$ with the ket-vectors $|\Psi\rangle$ in an abstract vector space. Not surprisingly, we can also define bra vectors. The bra vector $\langle\Psi|$ corresponding to the ket $|\Psi\rangle$ is identified with the complex conjugate function $\psi^*(x)$.

To use this idea effectively, we'll need to generalize some of the items in our mathematical tool kit. In earlier lectures, the labels that identified wave functions were members of some finite discrete set—for example, the eigenvalues of some observable. But now the independent variable is *continuous*. Among other things, this means that we cannot sum over it using ordinary sums. I think you know what to do, though. Here are function-oriented replacements for three of our vector-based concepts, two of which you will easily recognize:

- Integrals replace sums.

- Probability densities replace probabilities.
- Dirac delta functions replace Kronecker deltas.

Let's look at these items more closely.

Integrals Replace Sums: If we really wanted to be rigorous, we would begin by replacing the x axis by a discrete set of points separated by a very small distance ϵ , and then take the limit $\epsilon \rightarrow 0$. It would take several pages to justify each step. But we can avoid this trouble by a few intuitive definitions, such as replacing sums with integrals. Schematically, this concept can be written as

$$\sum_i \rightarrow \int dx.$$

For example, if we want to compute the area under a curve, we divide the x axis up into tiny segments and then add up the areas of a large number of rectangles, exactly as we do in elementary calculus. When we let the segments shrink to zero size, the sum becomes an integral.

Let's consider a bra $\langle \Psi |$ and a ket $|\Phi \rangle$ and define their inner product. The obvious way to do this is to replace the summation in Eq. 1.2 with an integral. We define the inner product to be

$$\langle \Psi | \Phi \rangle = \int_{-\infty}^{\infty} \psi^*(x) \phi(x) dx. \quad (8.2)$$

Probability Densities Replace Probabilities: Later, we will identify

$$P(x) = \psi^*(x)\psi(x)$$

as a probability density for the variable x . Why a probability density and not just a probability? If x is a continuous variable, then the probability that it will have any exact value is typically zero. A more useful question to ask is: What is the probability that x lies between two values, $x = a$ and $x = b$? Probability densities are defined so that this probability is given by an integral:

$$P(a, b) = \int_a^b P(x) dx = \int_a^b \psi^*(x)\psi(x) dx.$$

Because the total probability should be 1, we can define a normalized vector by

$$\int_{-\infty}^{\infty} \psi^*(x)\psi(x) dx = 1. \quad (8.3)$$

Dirac Delta Functions Replace Kronecker Deltas: So far, this should be very familiar. The Dirac delta function may be less so. The delta function is the analog of the Kronecker delta, δ_{ij} . The Kronecker delta is defined to be 0 for $i \neq j$ and 1 for $i = j$. But it can also be defined another way. Consider any vector F_i in a finite dimensional space. It is easy to see that the Kronecker delta satisfies

$$\sum_j \delta_{ij} F_j = F_i.$$

That's because the only nonzero term in the sum is the one where $j = i$. Within the summation, the Kronecker symbol filters out all the F 's except F_i . The obvious generalization is to define a new function that has similar filtering properties when used inside an integral. In other words, we want a new entity

$$\delta(x - x')$$

with the property that, for any function $F(x)$,

$$\int_{-\infty}^{\infty} \delta(x - x') F(x') dx' = F(x). \quad (8.4)$$

Eq. 8.4 defines this new entity, called the *Dirac delta function*, which turns out to be an essential tool in quantum mechanics. But despite its name, it isn't really a function in the usual sense. It is zero whenever $x \neq x'$, but when $x = x'$ it is infinite. In fact it is just infinite enough that the area under $\delta(x)$ equals 1. Roughly speaking, it is a function that is nonzero over an infinitesimal interval ϵ , but on that interval it has the value $1/\epsilon$. Thus, its area is 1, and, more importantly, it satisfies Eq. 8.4. The function

$$\frac{n}{\sqrt{\pi}} e^{-(nx)^2}$$

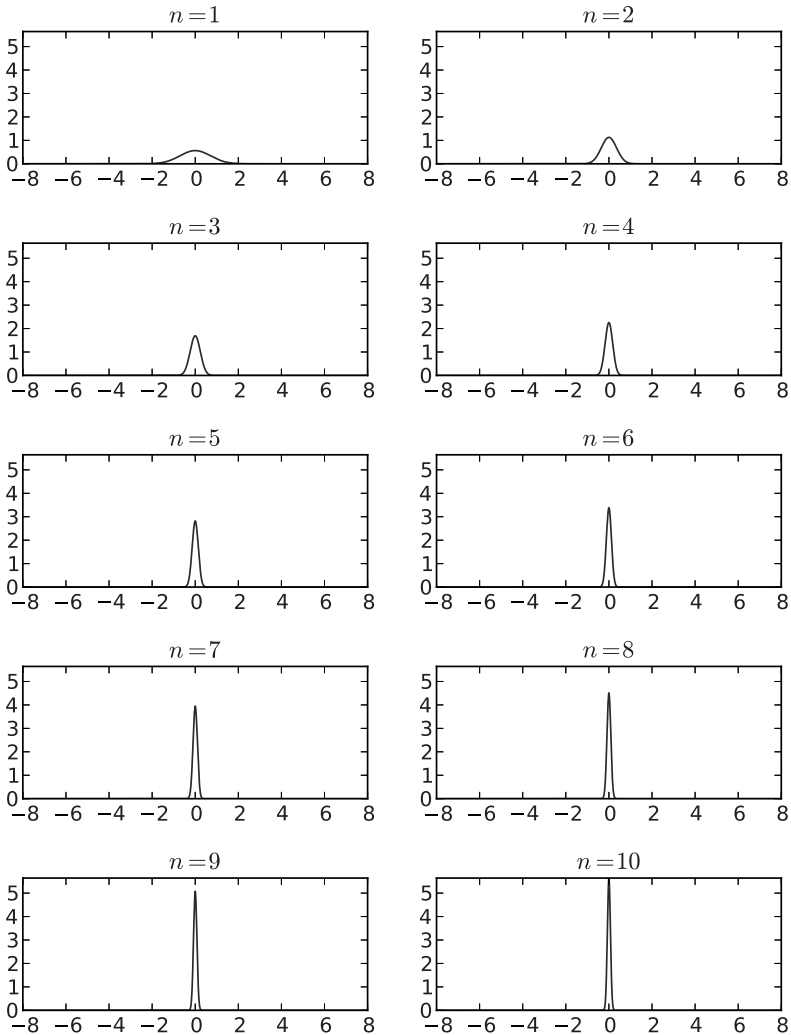


Figure 8.1: Dirac Delta Function Approximations. These approximations are based on $\frac{n}{\sqrt{\pi}}e^{-(nx)^2}$ and plotted for increasing values of n .

approximates the delta function reasonably well as n becomes very large. Fig. 8.1 plots this approximation for increasing values of n . Even though we stop at $n = 10$, a very small value, notice that the graph has already become very narrow and sharply peaked.

8.1.3 Integration by Parts

Before discussing linear operators, we'll take a short detour to remind you of a technique called *integration by parts*. It's fairly simple, and indispensable for our purposes. We'll be using it again and again. Suppose we take two functions, F and G , and consider the differential of their product FG . We can write

$$d(FG) = FdG + GdF$$

or

$$d(FG) - GdF = FdG.$$

Taking the definite integral gives us

$$\int_a^b d(FG) - \int_a^b GdF = \int_a^b FdG$$

or

$$FG \Big|_a^b - \int_a^b GdF = \int_a^b FdG.$$

This is the standard formula that you may remember from calculus. But in quantum mechanics the limits of integration tend to span the entire axis, and our wave functions must go to zero at infinity to be properly normalized. Therefore, the first term of this expression will always evaluate to zero. With that in mind, we can use a simplified version of integration by parts:

$$\int_{-\infty}^{\infty} F \frac{dG}{dx} dx = - \int_{-\infty}^{\infty} \frac{dF}{dx} G dx.$$

This form is correct as long as F and G go to zero appropriately at infinity, so that the boundary term becomes zero. You will do yourself a big favor if you just memorize this pattern: Switch the derivative from one factor of the integrand to the other at the cost of a minus sign.

8.1.4 Linear Operators

Bras and kets are half the story in quantum mechanics; the other half is the concept of linear operators and, in particular, Hermitian operators. This raises two questions:

- What is meant by a linear operator on a space of functions?
- What is the condition for a linear operator to be Hermitian?

The concept of a linear operator is simple enough: it's a machine that acts on a function and gives another function.

When it acts on the sum of two functions, it gives the sum of the individual results. When it acts on a complex numerical multiple of a function, it gives the same multiple of the original result. In other words, it is (surprise!) linear.

Let's look at some examples. One simple operation we can perform on a function $\psi(x)$ is to multiply it by x . That gives a new function $x\psi(x)$, and you can easily check that the action is linear. We'll represent the "multiply by x " operator with the symbol \mathbf{X} . By definition, then,

$$\mathbf{X} \psi(x) = x\psi(x). \quad (8.5)$$

Here's another example. Define \mathbf{D} to be the differentiation operator:

$$\mathbf{D} \psi(x) = \frac{d\psi(x)}{dx}. \quad (8.6)$$

Exercise 8.1: Prove that \mathbf{X} and \mathbf{D} are linear operators.

This, of course, is a minute subset of the possible linear operators that can be constructed, but we will soon see that \mathbf{X} and \mathbf{D} play a very central role in the quantum mechanics of particles.

Now, let's consider the property of Hermiticity. A convenient way to define a Hermitian operator is through its matrix elements, by sandwiching it between a bra and a ket. You can sandwich an operator \mathbf{L} in two different ways:

$$\langle \Psi | \mathbf{L} | \Phi \rangle$$

or

$$\langle \Phi | \mathbf{L} | \Psi \rangle.$$

In general, there is no simple relation between these two sandwiches. But in the case of a Hermitian operator (for which, by definition, $\mathbf{L}^\dagger = \mathbf{L}$) there is a simple relation: the two sandwiches are complex conjugates of each other:

$$\langle \Psi | \mathbf{L} | \Phi \rangle = \langle \Phi | \mathbf{L} | \Psi \rangle^*.$$

Let's see whether the operators \mathbf{X} and \mathbf{D} are Hermitian. Recalling that

$$\mathbf{X} \psi(x) = x\psi(x),$$

and using the inner product formula Eq. 8.2, we can write

$$\langle \Psi | \mathbf{X} | \Phi \rangle = \int \psi^*(x) x \phi(x) dx$$

and

$$\langle \Phi | \mathbf{X} | \Psi \rangle = \int \phi^*(x) x \psi(x) dx.$$

Because x is real, it's easy to see that these two integrals are complex conjugates of each other, and therefore that \mathbf{X} is Hermitian.

What about the operator \mathbf{D} ? In this case, the two sandwiches are

$$\langle \Psi | \mathbf{D} | \Phi \rangle = \int \psi^*(x) \frac{d\phi(x)}{dx} dx \quad (8.7)$$

and

$$\langle \Phi | \mathbf{D} | \Psi \rangle = \int \phi^*(x) \frac{d\psi(x)}{dx} dx. \quad (8.8)$$

To determine if \mathbf{D} is Hermitian, we need to compare these two integrals and see if they are complex conjugates of each other. In this form, it's a bit difficult to tell. The trick is to do the second integral by parts. As we explained, integration by parts allows you to switch the derivative from one factor in the integrand to the other, as long as you change the sign at the same time. Therefore, the integral in Eq. 8.8 can be rewritten as

$$\langle \Phi | \mathbf{D} | \Psi \rangle = - \int \psi(x) \frac{d\phi^*(x)}{dx} dx. \quad (8.9)$$

Now, we just need to compare the two expressions in Eqs. 8.7 and 8.9, which turns out to be easy. Because of the minus sign, it's clear that they are definitely *not* complex conjugates of each other. Instead, their relationship is captured by

$$\langle \Psi | \mathbf{D} | \Phi \rangle = - \langle \Phi | \mathbf{D} | \Psi \rangle^*,$$

which is the diametric opposite of what we wanted. Unlike the \mathbf{X} operator, \mathbf{D} is not Hermitian. Instead, it satisfies

$$\mathbf{D}^\dagger = -\mathbf{D}.$$

An operator with this property is called *anti-Hermitian*.

Although anti-Hermitian and Hermitian operators are opposites, it's very easy to go from one to the other. All you have to do is multiply by the imaginary number i or $-i$. Therefore we can use \mathbf{D} to construct an operator that *is* Hermitian, namely

$$-i\hbar\mathbf{D}.$$

If we look at the action of this new Hermitian operator on wave functions, we find that

$$-i\hbar\mathbf{D}\psi(x) = -i\hbar\frac{d\psi(x)}{dx}. \quad (8.10)$$

Keep this formula in mind. It will soon play a leading role in defining a very important property of particles—their momentum.

8.2 The State of a Particle

In classical mechanics, the “state of a system” means everything you need to know to predict the system’s future, given the forces acting on it. That, of course, means the positions of all the particles comprising the system, as well as the momenta of those particles. From a classical perspective, the instantaneous positions and momenta are entirely independent variables. For example, for a particle of mass m moving

along a one-dimensional axis x , the momentary state of the system is described by the pair (x, p) . The coordinate x is the location of the particle, and $p = m\dot{x}$ is its momentum. Taken together, these two variables define the phase space of the system. If we also know the force on the particle as a function of its position, Hamilton's equations permit us to calculate its position and momentum at all later times. They define a flow through the phase space.

Given this, one might guess that the quantum state of a particle would be spanned by a basis of states labeled by position and momentum:

$$|x, p\rangle.$$

The wave function would then be a function of both variables:

$$\psi(x, p) = \langle x, p | \Psi \rangle.$$

However, this is incorrect. We've already seen that things that would be simultaneously knowable in classical physics may not be in quantum mechanics. Different components of a spin, say σ_z and σ_x , are an example. One cannot know both components simultaneously; therefore, one does not have states in which both components are specified. The same is true for x and p : specifying both values is too much. Whether we're talking about spins (σ_z, σ_x) or positions and momenta (x, p) , the incompatibility is ultimately an experimental fact.

What then can we know about the particle on the x axis, if not x and p ? The answer is x **or** p ; for according to the mathematics of position and momentum operators, the two do not commute. But I emphasize that this is not something you could have predicted in advance; it is the distillation of many decades of experimental observations.

If the position of a particle is an observable, there must be a Hermitian operator associated with it. The obvious candidate is the operator \mathbf{X} . The first step in understanding this fundamental connection between the intuitive concept of position and the mathematical operator \mathbf{X} is to work out the eigenvectors and eigenvalues of \mathbf{X} . The eigenvalues are the possible values of position that can be observed, and the eigenvectors represent the states of definite position.

8.2.1 The Eigenvalues and Eigenvectors of Position

The obvious next question is: What are the possible outcomes of measuring \mathbf{X} , and what are the states in which it has a definite (predictable) value? In other words, what are its eigenvalues and eigenvectors? We'll start with \mathbf{X} . The eigen-equation for \mathbf{X} is

$$\mathbf{X}|\Psi\rangle = x_0|\Psi\rangle,$$

where the eigenvalue is denoted by x_0 . In terms of wave functions, this becomes

$$x\psi(x) = x_0\psi(x). \tag{8.11}$$

This last equation seems strange. How can x times a function be proportional to the same function? On the face of it, this seems impossible. But let's pursue it. We can rewrite Eq. 8.11 in the form

$$(x - x_0)\psi(x) = 0.$$

Of course, if a product is zero, then at least one of the factors must be zero. But the other factors may be different from zero. Thus, if $x \neq x_0$, then $\psi(x) = 0$. That's a very strong condition. It says that for a given eigenvalue x_0 , the function $\psi(x)$ can be nonzero at only one point, namely at

$$x = x_0.$$

For an ordinary continuous function this condition would be deadly: no sensible function can be zero everywhere except at one point, and be nonzero only at that point. But that is exactly the property of the Dirac delta function

$$\delta(x - x_0).$$

Evidently, then, every real number x_0 is an eigenvalue of \mathbf{X} , and the corresponding eigenvectors are functions (we often call them *eigenfunctions*) that are infinitely concentrated at $x = x_0$. The meaning of this is clear: the wave functions

$$\psi(x) = \delta(x - x_0)$$

represent states in which the particle is located right at the point x_0 on the x axis.

It of course makes a lot of sense that the wave function representing a particle known to be at x_0 is zero everywhere except at x_0 . How could it be otherwise? But it is gratifying to see the mathematics confirm this intuition.

Consider the inner product of a state $|\Psi\rangle$ and a position eigenstate $|x_0\rangle$:

$$\langle x_0|\Psi\rangle.$$

Using Eq. 8.2, we get

$$\langle x_0|\Psi\rangle = \int_{-\infty}^{\infty} \delta(x - x_0)\psi(x).$$

By the definition of delta functions given in Eq. 8.4, this integral evaluates to

$$\langle x_0|\Psi\rangle = \psi(x_0). \quad (8.12)$$

Because this is true for any x_0 , we can drop the subscript and write the general equation

$$\langle x|\Psi\rangle = \psi(x). \quad (8.13)$$

In other words, the wave function, $\psi(x)$, of a particle moving in the x direction is the projection of a state-vector $|\Psi\rangle$ onto the eigenvectors of position. We will also refer to $\psi(x)$ as the *wave function in the position representation*.

8.2.2 Momentum and Its Eigenvectors

Position is intuitive; momentum is less so, particularly in quantum mechanics. It will only be later that we see the connection between the operator that we identify with momentum and the familiar classical concept of mass times velocity. But I assure you that we will make the connection.

For now, let's take the abstract mathematical route. The momentum operator in quantum mechanics is called \mathbf{P} , and it is defined in terms of the operator $-i\mathbf{D}$:

$$-i\mathbf{D} = -i \frac{d}{dx}.$$

As we saw earlier in Eq. 8.10, we need the factor $-i$ to make this operator Hermitian.

We could just define \mathbf{P} to be $-i\mathbf{D}$, but if we did, we would run into a problem later when we connect these ideas to those of classical physics. The reason should be clear—there's a dimensional mismatch. In classical physics, the units of momentum are mass times velocity—in other words, mass times length divided by time (ML/T). On the other hand, the operator \mathbf{D} has units of inverse length, or $1/L$. The resolution of the mismatch is provided by Planck's constant \hbar , which has units of ML^2/T . The correct relation between \mathbf{P} and \mathbf{D} is therefore

$$\mathbf{P} = -i\hbar\mathbf{D} \tag{8.14}$$

or, in terms of its action on wave functions,

$$\mathbf{P}\psi(x) = -i\hbar \frac{d\psi(x)}{dx}. \tag{8.15}$$

Quantum physicists often use units in which \hbar is exactly one, and in that way simplify the equations. As tempting as it is, we won't do that here.

Let's work out the eigenvectors and eigenvalues of \mathbf{P} . The eigen-equation in abstract vector notation is

$$\mathbf{P}|\Psi\rangle = p|\Psi\rangle, \quad (8.16)$$

where the symbol p is an eigenvalue of \mathbf{P} . Eq. 8.16 can also be expressed in terms of wave functions. Using the identification

$$\mathbf{P} = -i\hbar \frac{d}{dx},$$

we can write the eigen-equation as

$$-i\hbar \frac{d\psi(x)}{dx} = p\psi(x)$$

or

$$\frac{d\psi(x)}{dx} = \frac{ip}{\hbar}\psi(x).$$

This is a type of equation that we've run into before. The solution has the form of an exponential:

$$\psi_p(x) = Ae^{\frac{ipx}{\hbar}}.$$

The subscript p is just a reminder that $\psi_p(x)$ is the eigenvector of \mathbf{P} with the specific eigenvalue p . It is a function of x , but it is labeled by an eigenvalue of \mathbf{P} .

The constant A multiplying the exponential is not determined by the eigenvector equation. That's nothing new; the eigenvalue equation never tells us the overall normalization of the wave function. As a rule, we fix the constant by requiring the wave function to be normalized to unit probability. An example that goes all the way back to Section 2.3 is the eigenvector of the x component of spin:

$$|r\rangle = \frac{1}{\sqrt{2}}|u\rangle + \frac{1}{\sqrt{2}}|d\rangle.$$

The factor $1/\sqrt{2}$ is there to make sure the total probability is 1.

Normalizing the eigenvectors of \mathbf{P} is a more subtle operation, but the result is simple. The factor A is only slightly more complicated than in the spin case. To save time, I will tell you the answer and leave it for you to prove later. The correct factor is $A = 1/\sqrt{2\pi}$. Thus,

$$\psi_p(x) = \frac{1}{\sqrt{2\pi}} e^{\frac{ipx}{\hbar}}. \quad (8.17)$$

A point of some interest follows from Eqs. 8.13 and 8.17. The inner product of a position eigenvector $|x\rangle$ and a momentum eigenvector $|p\rangle$ has a very simple and symmetric form:

$$\begin{aligned} \langle x|p\rangle &= \frac{1}{\sqrt{2\pi}} e^{\frac{ipx}{\hbar}} \\ \langle p|x\rangle &= \frac{1}{\sqrt{2\pi}} e^{\frac{-ipx}{\hbar}}. \end{aligned} \quad (8.18)$$

The second equation is simply the complex conjugate of the first. These results are easy to verify if you keep in mind that $|x\rangle$ is represented by a delta function. I'd like to mention two important points before moving further:

1. Eq. 8.17 represents a momentum eigenfunction in the position basis. In other words, although it represents a momentum eigenstate, it is a function of x , and *not* an explicit function of p .
2. We've been using the symbol ψ for both position and momentum eigenstates. A mathematician might not approve of using the same symbol for two different functions, but physicists do it all the time. $\psi(x)$ is just the generic symbol for whatever function we happen to be discussing.

At this juncture, we begin to get a glimmer of why the wave function is called the wave function. What you should notice is that the eigenfunctions (wave functions representing eigenvectors) of the momentum operator have the form of waves—sine waves and cosine waves, to be precise. In fact, we can now see one of the most fundamental aspects of the wave-particle duality of quantum mechanics. The wavelength of the function

$$e^{\frac{ipx}{\hbar}}$$

is given by

$$\lambda = \frac{2\pi\hbar}{p}$$

because the value of the function is unchanged if we add $\frac{2\pi\hbar}{p}$ to the variable x :

$$e^{\frac{ip(x + \frac{2\pi\hbar}{p})}{\hbar}} = e^{\frac{ipx}{\hbar}} e^{2\pi i} = e^{\frac{ipx}{\hbar}}.$$

Let's pause for a moment to discuss the importance of this connection between momentum and wavelength. It's not just important: in many ways, it is the relationship that defined twentieth-century physics. Over the last hundred years, physicists have primarily been concerned with uncovering the laws of the microscopic world. This has meant figuring out how objects are built out of smaller objects. The examples are obvious: molecules are made from atoms; atoms from electrons and nuclei; nuclei from protons and neutrons. These subnuclear particles are constructed out of quarks and gluons. And the game goes on as scientists search for ever smaller and more hidden entities.

All of these objects are too small to see with the best optical microscopes, let alone the naked eye. The reason is not just that our eyes are insufficiently sensitive. The more important fact is that eyes and optical microscopes are sensitive to the visible spectrum, which comprises wavelengths at least a few thousand times longer than the size of an atom. As a rule, you can't resolve objects much smaller than the wavelength you're using to look at them. For this reason, the story of twentieth-century physics was in large part a quest

for smaller and smaller wavelengths of light—or any other kind of wave. In Lecture 10, we will discover that light of a given wavelength is composed of photons whose momentum is related to the wavelength by exactly the relation

$$\lambda = \frac{2\pi\hbar}{p}.$$

The implication is that to probe objects of ever smaller size one needs photons (or other objects) of ever larger momentum. Large momentum inevitably means large energy. It's for that reason that the discovery of the microscopic properties of matter required increasingly powerful particle accelerators.

8.3 Fourier Transforms and the Momentum Basis

The wave function $\psi(x)$ has the important role of determining the probability for finding the particle at position x :

$$P(x) = \psi^*(x)\psi(x).$$

As we will see, no experiment can determine both the position and momentum of a particle simultaneously. But if we forego determining anything about the position, momentum can be measured precisely. The situation is quite analogous to that of the x and z components of a spin. Either value can be measured, but not both.

What is the probability that a particle has momentum p if we choose to measure it? The answer is a straightforward generalization of the principles laid down in Lecture 3. The probability that a momentum measurement will give momentum p is

$$P(p) = | \langle \mathbf{P} | \Psi \rangle |^2. \quad (8.19)$$

The entity $\langle \mathbf{P} | \Psi \rangle$ is called the wave function of $|\Psi\rangle$ in the *momentum representation*. Naturally, it is a function of p and is denoted by a new symbol:

$$\tilde{\psi}(p) = \langle \mathbf{P} | \Psi \rangle. \quad (8.20)$$

It is now clear that there are two ways to represent a state-vector. One way is in the position basis and the other is in the momentum basis. Both wave functions—the position wave function $\psi(x)$ and the momentum wave function $\tilde{\psi}(p)$ —represent exactly the same state-vector $|\Psi\rangle$. It follows that there must be some transformation between them such that if you know $\psi(x)$, the transformation produces $\tilde{\psi}(p)$, and vice versa. In fact, the two representations are Fourier transforms of each other.

8.3.1 Resolving the Identity

We are about to see the great power of the Dirac bra-ket notation in simplifying complicated things. First, let's recall an important idea from earlier lectures. Suppose we define an orthonormal basis of states through the eigenvectors of

some Hermitian observable. Call the basis vectors $|i\rangle$. In Lecture 7, I explained a very useful trick, and now we are going to see just how useful it is. It's called *resolving the identity*. The trick given in (Eq. 7.11) is to write the identity operator \mathbf{I} (the operator that acts on any vector to give the same vector) in the form

$$\mathbf{I} = \sum_i |i\rangle\langle i|.$$

Because momentum and position are both Hermitian, the sets of vectors $|x\rangle$ and set $|p\rangle$ each define basis vectors. By replacing summation with integration we discover two ways to resolve the identity:

$$\mathbf{I} = \int dx |x\rangle\langle x| \quad (8.21)$$

and

$$\mathbf{I} = \int dp |p\rangle\langle p|. \quad (8.22)$$

Let's suppose that we know the wave function of the abstract vector $|\Psi\rangle$ in the position representation. By definition, it is equal to

$$\psi(x) = \langle x|\Psi\rangle. \quad (8.23)$$

Now suppose we want to know the wave function $\tilde{\psi}(p)$ in the momentum representation. Here are the steps laid out in detail:

- First, use the definition of the momentum-representation wave function:

$$\tilde{\psi}(p) = \langle p | \Psi \rangle.$$

- Now, insert the unit operator between the bra- and ket-vectors, in the form given in Eq. 8.21:

$$\tilde{\psi}(p) = \int dx \langle p | x \rangle \langle x | \Psi \rangle.$$

- The expression $\langle x | \Psi \rangle$ is just the wave function $\psi(x)$, and $\langle p | x \rangle$ is given to us by the second equation of Eqs. 8.18:

$$\langle p | x \rangle = \frac{1}{\sqrt{2\pi}} e^{\frac{-ipx}{\hbar}}.$$

- Putting it all together, we find that

$$\tilde{\psi}(p) = \frac{1}{\sqrt{2\pi}} \int dx e^{\frac{-ipx}{\hbar}} \psi(x). \quad (8.24)$$

This equation shows us exactly how to transform a given wave function in the position representation into the corresponding wave function in the momentum representation. What is it good for? Suppose the position wave function for some particle is known; however, the goal of your experiment is to measure the momentum, and you want to know

the probability of observing momentum p . The procedure is to first calculate $\tilde{\psi}(p)$ by using Eq. 8.24 and then compute the probability

$$P(p) = \tilde{\psi}^*(p)\tilde{\psi}(p).$$

It's just as easy to go the other way. Suppose we know $\tilde{\psi}(p)$ and wish to recover $\psi(x)$. This time, we use Eq. 8.22 to resolve the identity. Here are the steps (notice that they look suspiciously similar to the earlier ones):

- First, use the definition of the position-representation wave function:

$$\psi(x) = \langle x|\Psi\rangle$$

- Now, insert the unit operator between the bra- and ket-vectors, in the form given in Eq. 8.22:

$$\psi(x) = \int dp \langle x|p\rangle \langle p|\Psi\rangle.$$

- The expression $\langle p|\Psi\rangle$ is just the wave function $\tilde{\psi}(p)$, and $\langle x|p\rangle$ is given to us by Eqs. 8.18. But this time, it's the first of the two equations.

$$\langle x|p\rangle = \frac{1}{\sqrt{2\pi}} e^{\frac{ipx}{\hbar}}.$$

- Putting it all together, we find that

$$\psi(x) = \frac{1}{\sqrt{2\pi}} \int dp e^{\frac{ipx}{\hbar}} \tilde{\psi}(p).$$

Let's take another look at the two equations for going back and forth from position to momentum. Notice how symmetrical they are. The only asymmetry is that one equation contains $e^{\frac{ipx}{\hbar}}$ and the other contains $e^{\frac{-ipx}{\hbar}}$:

$$\begin{aligned} \tilde{\psi}(p) &= \frac{1}{\sqrt{2\pi}} \int dx e^{\frac{-ipx}{\hbar}} \psi(x) \\ \psi(x) &= \frac{1}{\sqrt{2\pi}} \int dp e^{\frac{ipx}{\hbar}} \tilde{\psi}(p). \end{aligned} \quad (8.25)$$

The relation between the position and momentum representations summarized by Eqs. 8.25 is that they are reciprocal *Fourier transforms* of one another. In fact, these are the central equations in the field of Fourier analysis. I want you to notice how easy it was to derive those equations using Dirac's elegant notation.

8.4 Commutators and Poisson Brackets

Earlier, in Lecture 4, we formulated two important principles about commutators. The first had to do with the connection between classical mechanics and quantum mechanics;

the second had to do with uncertainty. I now will finish up this very long lecture by showing you what these principles have to do with \mathbf{X} and \mathbf{P} .

We'll start with the connection between commutators and classical physics. As you may recall, we found that commutators have a great similarity to Poisson brackets, a relationship we made explicit in Eq. 4.21. If we plug in the operator symbols \mathbf{L} and \mathbf{M} that we've been using in this lecture, we get

$$[\mathbf{L}, \mathbf{M}] \iff i\hbar\{L, M\}, \quad (8.26)$$

and we're reminded that the equations for quantum motion strongly resemble their classical equivalents. This suggests that we may learn something by computing the commutator of the observables \mathbf{X} and \mathbf{P} . Fortunately, this is easy to do.

First, let's see what the product \mathbf{XP} does when it acts as an operator on an arbitrary wave function $\psi(x)$. Recalling Eqs. 8.5 and 8.15, we can write

$$\mathbf{X}\psi(x) = x\psi(x)$$

$$\mathbf{P}\psi(x) = -i\hbar\frac{d\psi(x)}{dx}.$$

Together, these equations tell us how the product \mathbf{XP} acts on $\psi(x)$:

$$\mathbf{XP}\psi(x) = -i\hbar x\frac{d\psi(x)}{dx} \quad (8.27)$$

Now, let's try it with \mathbf{X} and \mathbf{P} in the opposite order:

$$\mathbf{P}\mathbf{X}\psi(x) = -i\hbar \frac{d(x\psi(x))}{dx}.$$

To calculate this last expression, we just use the standard rule for differentiating the product $x\psi(x)$. Using this rule, it's easy to see that

$$\mathbf{P}\mathbf{X}\psi(x) = -i\hbar x \frac{d\psi(x)}{dx} - i\hbar\psi(x). \quad (8.28)$$

Now, we'll subtract Eq. 8.28 from Eq. 8.27 to show how the commutator acts on the wave function:

$$[\mathbf{X}, \mathbf{P}]\psi(x) = \mathbf{X}\mathbf{P}\psi(x) - \mathbf{P}\mathbf{X}\psi(x)$$

or

$$[\mathbf{X}, \mathbf{P}]\psi(x) = i\hbar\psi(x).$$

In other words, when the commutator $[\mathbf{X}, \mathbf{P}]$ acts on *any* wave function $\psi(x)$, all it does is multiply $\psi(x)$ by the number $i\hbar$. We can express this by writing

$$[\mathbf{X}, \mathbf{P}] = i\hbar. \quad (8.29)$$

This in itself is extremely important. The fact that \mathbf{X} and \mathbf{P} don't commute is the key to understanding why they are not simultaneously measurable. But things get even more interesting when we compare this equation with Equivalence 8.26, which relates commutators to classical Poisson brackets. In

fact, Eq. 8.29 suggests that the corresponding classical Poisson bracket is

$$\{x, p\} = 1,$$

which is exactly the classical relation between coordinates and their conjugate momenta (see *Volume I*, Lecture 10, Eq. 8). Ultimately, it is this connection that explains why the quantum concept of momentum is connected to the classical concept.

Using the general uncertainty principle from Lecture 5, we can now specialize to the case

$$[\mathbf{X}, \mathbf{P}] = i\hbar.$$

and

$$\Delta\mathbf{X}\Delta\mathbf{P} \geq \frac{\hbar}{2}.$$

We'll do that in the next section.

Now let's recall the second principle involving commutators. In Lecture 4, we found that two observables \mathbf{L} and \mathbf{M} cannot be determined simultaneously unless they commute. If they don't commute, you cannot measure \mathbf{L} without interfering with a measurement of \mathbf{M} . It is not possible to find simultaneous eigenvectors of two noncommuting observables. This led to the general uncertainty principle.

8.5 The Heisenberg Uncertainty Principle

And now, ladies and gentlemen, here's what you've all been waiting for. At long last: *the Heisenberg Uncertainty Principle*.

The Heisenberg Uncertainty Principle is one of the most famous results of quantum mechanics: it not only asserts that the position and momentum of a particle cannot be simultaneously known, but it also provides an exact quantitative limit for their mutual uncertainties. At this point, I suggest that you revisit Lecture 5, where I explained the general uncertainty principle. We did all the work there, and now we get to reap the benefits.

As we've seen, the general uncertainty principle puts a quantitative limit on the simultaneous uncertainties of two observables \mathbf{A} and \mathbf{B} . This idea was captured in Inequality 5.13:

$$\Delta\mathbf{A} \Delta\mathbf{B} \geq \frac{1}{2} |\langle \Psi | [\mathbf{A}, \mathbf{B}] | \Psi \rangle|.$$

Now let's apply this principle directly to the position and momentum operators \mathbf{X} and \mathbf{P} . In this case, the commutator is just a number and its expectation is that same number. Replacing \mathbf{A} and \mathbf{B} with \mathbf{X} and \mathbf{P} gives

$$\Delta\mathbf{X} \Delta\mathbf{P} \geq \frac{1}{2} |\langle \Psi | [\mathbf{X}, \mathbf{P}] | \Psi \rangle|,$$

and replacing $[\mathbf{X}, \mathbf{P}]$ with $i\hbar$ results in

$$\Delta \mathbf{X} \Delta \mathbf{P} \geq \frac{1}{2} |i\hbar \langle \Psi | \Psi \rangle|.$$

But $\langle \Psi | \Psi \rangle$ equals 1, and the end result is

$$\Delta \mathbf{X} \Delta \mathbf{P} \geq \frac{1}{2} \hbar.$$

No experiment can ever beat this limitation. You can try your best to determine a particle's momentum and position simultaneously in a reproducible manner, but no matter how careful you are, the uncertainty in the position times the uncertainty in the momentum will never be less than $\frac{1}{2}\hbar$.

As we saw in Section 8.2.1, the wave function of an eigenstate of \mathbf{X} is highly concentrated about some point x_0 ; in this eigenstate, the probability is also perfectly localized. On the other hand, the probability $P(x)$ for a momentum eigenstate is uniformly spread over the entire x axis. To see this, let's take the wave function in Eq. 8.17 and multiply it by its complex conjugate:

$$\psi_p^*(x)\psi_p(x) = \left(\frac{1}{\sqrt{2\pi}}e^{\frac{-ipx}{\hbar}}\right)\left(\frac{1}{\sqrt{2\pi}}e^{\frac{ipx}{\hbar}}\right) = \frac{1}{2\pi}.$$

The result is completely uniform, with no peaks anywhere on the x axis. Evidently, a state with definite momentum is completely uncertain in its position.

Fig. 8.2 illustrates the definition of uncertainty for the position variable x . In the top half of the figure, you can see that the uncertainty Δx is a measure of how spread out the function is in relation to its expectation value $\langle x \rangle$. The

label d shows the deviation of one point in relation to $\langle x \rangle$; this may be a positive or negative quantity. The uncertainty Δx is the result of an averaging process over all possible d 's and characterizes the function as a whole. To prevent the positive d 's from canceling the negative ones, each d value is squared during this averaging process.

The bottom half of Fig. 8.2 shows how the calculation can be simplified by shifting the origin to coincide with $\langle x \rangle$. The numerical value of Δx is unchanged by this shift.

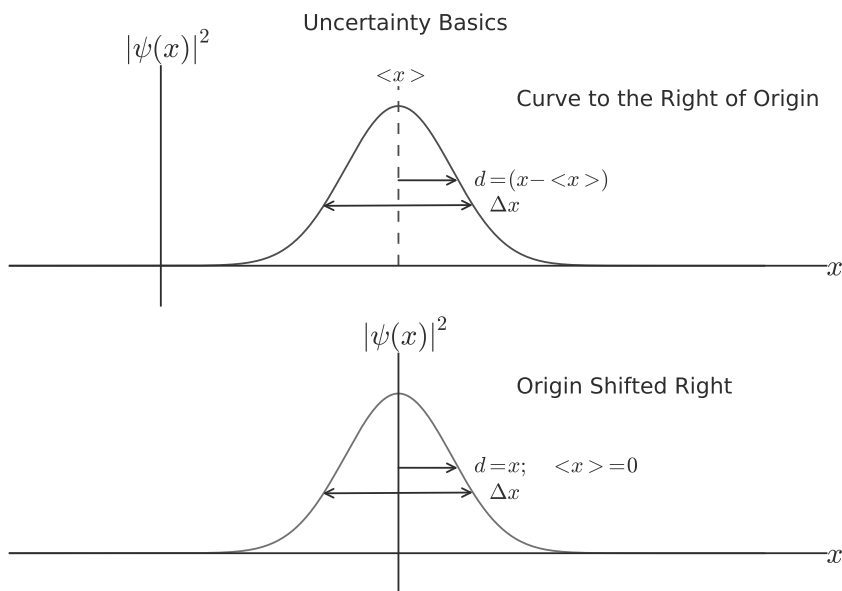


Figure 8.2: Uncertainty Basics. Top: $\langle x \rangle$ to right of origin. Deviations d may be positive or negative. Overall uncertainty Δx (> 0) derived from the average value of d^2 . Bottom: Origin shifted right, $\langle x \rangle = 0$, Δx has same value.

Lecture 9

Particle Dynamics

Art and Lenny expected some action at Hilbert's Place. But all the state-vectors were absolutely still—frozen, you might say.

Lenny: This is boring, Art. Doesn't anything ever happen around here? Hey Hilbert, why is this joint so still?

Hilbert: Oh, don't worry. Things will pick up as soon as the Hamiltonian gets here.

Art: The Hamiltonian? He sounds like a real operator.

9.1 A Simple Example

The first two volumes of the Theoretical Minimum series have largely focused on two questions. The first is: What do we mean by a system and how do we describe the momentary states of a system? As we've seen, the classical and quantum answers to this question are very different. Classical phase

space—the space of coordinates and momenta—is replaced in quantum theory by the linear vector space of states.

The second big question is: How do states change with time? In both classical mechanics and quantum mechanics, the answer is *according to the minus first law*. In other words, states change so that information and distinctions are never erased. In classical mechanics, this principle led to Hamilton's equations and Liouville's theorem. Earlier, in Lecture 4, I explained how in quantum mechanics this law led to the principle of unitarity, which in turn led to the general Schrödinger equation.

Lecture 8 was all about the first question: How do we describe the state of a particle? Now, in the current lecture, we come to the second question, which we might rephrase: *How do particles move in quantum mechanics?*

In Lecture 4, I laid out the basic rules for how quantum states change with time. The essential ingredient is the Hamiltonian H , which in both classical and quantum mechanics represents the total energy of a system. In quantum mechanics, the Hamiltonian controls the time evolution of a system through the time-dependent Schrödinger equation:

$$i\hbar \frac{\partial |\Psi\rangle}{\partial t} = \mathbf{H}|\Psi\rangle. \quad (9.1)$$

This lecture is all about the *Original Schrödinger Equation*—the equation that Schrödinger wrote down to describe a quantum mechanical particle. The Original Schrödinger Equation is a special case of Eq. 9.1.

The motion of ordinary (nonrelativistic) particles in classical mechanics is governed by a Hamiltonian, equal to the

kinetic energy plus the potential energy. We will soon come to the quantum version of this Hamiltonian, but first let's look at a Hamiltonian that's even simpler.

We'll start with the simplest Hamiltonian I can think of. In this case, the Hamiltonian operator \mathbf{H} is a fixed constant times the momentum operator \mathbf{P} :

$$\mathbf{H} = c\mathbf{P}. \quad (9.2)$$

This example is rarely written down, though it turns out to be quite instructive. The constant c is a fixed number. Is $c\mathbf{P}$ a reasonable Hamiltonian for a particle? Yes it is, and in a moment we'll find out what kind of particle it describes. For now, just notice that Eq. 9.2 is different from what we might expect for a nonrelativistic particle. In other words, it's not $\mathbf{P}^2/2m$. This simpler example is worth exploring first, just to see how the mathematical apparatus works.

How do we represent this example in terms of wave functions $\psi(x)$ in the position basis? We'll start by plugging our operators into the time-dependent Schrödinger equation (Eq. 9.1):

$$i\hbar \frac{\partial\psi(x,t)}{\partial t} = -ci\hbar \frac{\partial\psi(x,t)}{\partial x}.$$

Notice that we're now writing ψ as a function of both x and t . Canceling the $i\hbar$ terms gives us

$$\frac{\partial\psi(x,t)}{\partial t} = -c \frac{\partial\psi(x,t)}{\partial x}, \quad (9.3)$$

which is a pretty simple equation. In fact, any function of $(x - ct)$ is a solution. By “function of $(x - ct)$,” I mean any function that depends not on x and t separately, but only on the combination $(x - ct)$. To see how this works, just consider an arbitrary function $\psi(x - ct)$ and look at its derivatives. If you take the partial with respect to x , you just get

$$\frac{\partial\psi(x - ct)}{x}$$

because the derivative of $(x - ct)$ with respect to x is 1. But if you take the partial with respect to t , you get

$$-c\frac{\partial\psi(x - ct)}{\partial t}.$$

It’s clear that this combination of derivatives satisfies Eq. 9.3; therefore any function of this form solves the Schrödinger equation.

Now, let’s see how a function $\psi(x - ct)$ behaves. What does it look like? How does it evolve with time? Suppose we start by looking at a snapshot at $t = 0$. We can call the snapshot $\psi(x)$ because it tells us what ψ looks like at every point in space *at the specific time* $t = 0$. Of course, we don’t want just any function of $(x - ct)$. We want the total probability

$$\int_{-\infty}^{\infty} \psi^*(x)\psi(x)dx$$

to equal 1. In other words, we want $\psi(x)$ to fall off nicely to zero at infinity so that the integral doesn’t blow up. Fig.

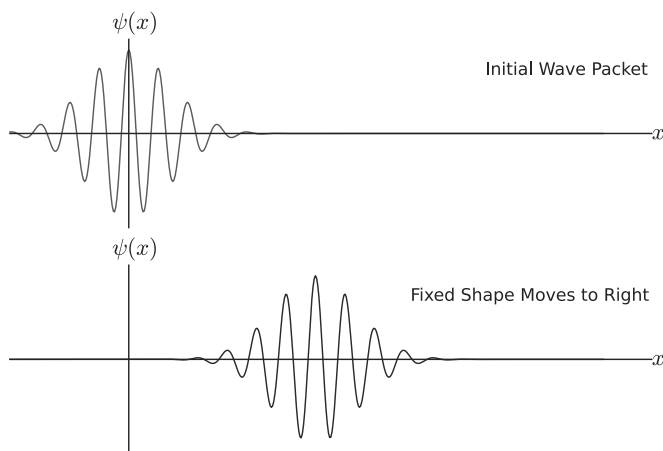


Figure 9.1: Fixed Shape Wave Packet Moving at Fixed Speed c

9.1 shows $\psi(x)$ schematically. With these characteristics, it makes sense to call $\psi(x)$ a wave *packet*.

Now that we've described the snapshot $\psi(x)$ at $t = 0$, what happens if we let time move forward? As t increases, the wave packet keeps the exact same shape. Every feature of the complex-valued function $\psi(x, t)$ moves with uniform velocity c to the right.¹

I had a reason for giving the name c to our constant—the symbol c often stands for the speed of light. So is this particle a photon? No, not really. But our description of this hypothetical particle is pretty close to the correct description of a neutrino that moves at the speed of light. (Real

¹This includes both the real and the imaginary parts of $\psi(x)$.

neutrinos probably move at a speed that is immeasurably smaller than the speed of light.) This Hamiltonian would be a very good description of a one-dimensional neutrino except for one problem: the particle described by our wave function can only move to the right. To round out this description, we would have to add another possibility—that the particle could also move to the left!²

Our right-going zaxon³ has another oddball feature—its energy can be either positive or negative. This is because the \mathbf{P} operator, as a vector, can take on positive or negative values. In general, the energy of a particle with negative momentum is negative, and the energy of a particle with positive momentum is positive. I won't say more about this except that the problem of negative energy for this kind of particle was solved by Dirac, who used it to establish the theoretical basis for antiparticles. For our purposes, we can ignore this issue and simply allow the energy of our particle to be either positive or negative.

Since the wave function of our particle moves rigidly down the x axis, so does the probability distribution. As a result, the expectation value of x moves in exactly the same way, which is to say that it moves to the right with velocity c . That's the essential quantum mechanics of this system. However, there is another important thing to keep in mind. When we said the velocity c is a *fixed* constant, we

²Our right-going particles remind me of Dr. Seuss's classic story "The Zax," and I'm tempted to call them "right-going zaxons." There's no telling how the story would have turned out if Theodor Geisel had known more about neutrinos.

³There. I've said it.

weren't kidding. Our particle can only exist in a state where it moves at this particular velocity. It can never slow down or speed up.

How does this compare with the classical description of such a particle? Starting with the same Hamiltonian, a classical physicist would just write Hamilton's equations. With $H = cP$, Hamilton's equations are

$$\frac{\partial H}{\partial p} = \dot{x}$$

and

$$\frac{\partial H}{\partial x} = -\dot{p}.$$

Carrying out the partial derivatives, these become

$$\frac{\partial H}{\partial p} = \dot{x} = c$$

and

$$\frac{\partial H}{\partial x} = -\dot{p} = 0.$$

Thus, in the classical description of our particle, the momentum is conserved, and the position moves with fixed velocity c . In the quantum mechanical description, the whole probability distribution and the expectation value move with velocity c . In other words, the expectation value of position behaves according to the classical equations of motion.

9.2 Nonrelativistic Free Particles

Only massless particles can move at the velocity of light, and I might add, they can only move at that velocity. All known particles other than photons and gravitons are massive and can move at any velocity less than c . When they move with a velocity *much* less than c , they are said to be nonrelativistic and their motion is governed by ordinary Newtonian mechanics, at least classically. The earliest application of quantum mechanics was to the motion of nonrelativistic particles.

I showed earlier (in Lectures 4 and 8) that Poisson brackets play the same mathematical role in classical mechanics as commutators do in quantum mechanics. Written with these constructs, the classical and quantum mechanical equations of motion are almost identical in form. In particular, the Hamiltonian comes into play in the same way with Poisson brackets as it does with commutators. So, if you want to write down the quantum mechanical equations of a system whose classical physics you already know, it's very reasonable to try using the classical Hamiltonian, translated into operator form.

For a nonrelativistic free particle, the natural Hamiltonian to try is $p^2/2m$. When we say the particle is free, what we really mean is that no forces are acting on it, and therefore we can ignore potential energy. All we care about is the kinetic energy, which is defined as

$$T = \frac{1}{2}mv^2.$$

As you recall, the momentum for a classical particle is

$$p = mv.$$

The Hamiltonian is just the kinetic energy, which we can write in terms of the momentum p . This gives us

$$H = \frac{1}{2}mv^2 = \frac{p^2}{2m}$$

for the Hamiltonian of a classical nonrelativistic free particle. Unlike the right-going zaxon of the previous example, the energy of this particle does not depend on its direction of motion. That's because the energy is proportional to p^2 rather than p itself. So we'll start with a particle whose energy is $p^2/2m$ and work out the Schrödinger equation (the original one that Schrödinger discovered) for a free particle.

Our plan is to follow the same process we used in the previous example, using the Hamiltonian to write a time-dependent Schrödinger equation. As usual, the left side of the equation is

$$i\hbar \frac{\partial \psi}{\partial t}.$$

We'll derive the right-hand side by rewriting the classical Hamiltonian—the kinetic energy—as an operator. The classical kinetic energy is

$$p^2/2m.$$

The quantum version replaces p with \mathbf{P} :

$$\mathbf{H} = \mathbf{P}^2/2m.$$

What is the meaning of this? As we've seen, the operator \mathbf{P} is defined as

$$\mathbf{P} = -i\hbar \frac{\partial}{\partial x}.$$

The square of \mathbf{P} is just the operator that you get by allowing \mathbf{P} to act twice in succession. Thus,

$$\mathbf{P}^2 = \left(-i\hbar \frac{\partial}{\partial x}\right)\left(-i\hbar \frac{\partial}{\partial x}\right),$$

or

$$\mathbf{P}^2 = -\hbar^2 \frac{\partial^2}{\partial x^2},$$

and the Hamiltonian becomes

$$\mathbf{H} = -\frac{\hbar^2}{2m} \frac{\partial^2}{\partial x^2}.$$

Finally, if we equate the left- and right-hand sides of the time-dependent Schrödinger equation, we get

$$i\hbar \frac{\partial \psi}{\partial t} = \frac{-\hbar^2}{2m} \frac{\partial^2 \psi}{\partial x^2}. \quad (9.4)$$

This is the traditional Schrödinger equation for an ordinary nonrelativistic free particle. It is a particular kind of wave equation, but, in contrast to the previous example, waves

of different wavelength (and momenta) move with different velocities. Because of this, the wave function *does not* maintain its shape. Unlike the zaxxon wave function, it tends to spread out and fall apart. This is shown schematically in Fig. 9.2.

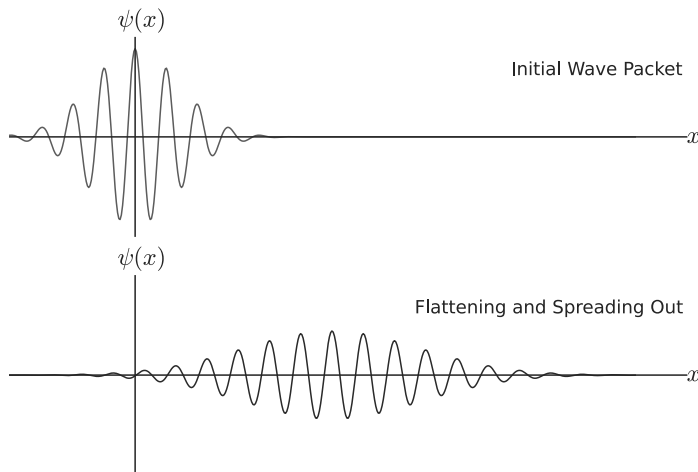


Figure 9.2: Typical Wave Packet for a Nonrelativistic Free Particle. Top: The initial wave packet is compact and highly localized. Bottom: Over time, the wave packet moves to the right and spreads out.

9.3 Time-Independent Schrödinger Equation

We are going to solve the time-dependent Schrödinger equation for nonrelativistic free particles, but first we need to

solve the time-independent version. The time-independent equation is essentially the eigenvector equation for the Hamiltonian,

$$H|\Psi\rangle = E|\Psi\rangle,$$

written explicitly in terms of the wave function $\psi(x)$:

$$-\frac{\hbar^2}{2m} \frac{\partial^2 \psi(x)}{\partial x^2} = E\psi(x). \quad (9.5)$$

It's very easy to find a complete set of eigenvectors that satisfy this equation. In fact, momentum eigenvectors do the job. Let's try the function

$$\psi(x) = e^{\frac{ipx}{\hbar}} \quad (9.6)$$

as a possible solution. Carrying out the derivatives, we find that this function is indeed a solution to Eq. 9.5, as long as we set

$$E = p^2/2m. \quad (9.7)$$

This should come as no surprise—after all, E represents an energy eigenvalue in Eq. 9.5.

Exercise 9.1: Derive Eq. 9.7 by plugging Eq. 9.6 into Eq. 9.5.

As we saw in Section 4.13, every solution to the time-independent Schrödinger equation allows us to construct to a time-dependent

solution. All we need to do is multiply the time-independent solution—in this case $e^{\frac{ipx}{\hbar}}$ —by $e^{-i\frac{Et}{\hbar}} = e^{-i\frac{p^2t}{2m\hbar}}$. Thus, a complete set of solutions can be written as

$$\psi(x, t) = \exp \frac{i(px - \frac{p^2t}{2m})}{\hbar}.$$

Any solution is a sum, or integral, of these solutions:

$$\psi(x, t) = \int \tilde{\psi}(p) \left(\exp \frac{i(px - \frac{p^2t}{2m})}{\hbar} \right) dp.$$

You can start with any wave function at $t = 0$, find $\tilde{\psi}(p)$ by Fourier transform, and let it evolve. The shape will change because the waves for different p values travel at different velocities. But, as we will soon see, the overall wave packet will travel at velocity $\langle p/m \rangle$ just as a classical particle would.

This simple general solution has an important implication. Among other things, it says that the momentum-representation wave function changes with time in a very simple way:

$$\tilde{\psi}(p, t) = \tilde{\psi}(p) \exp \frac{i(px - \frac{p^2t}{2m})}{\hbar}.$$

In other words, only the phase changes with time, while the magnitude remains constant. What makes this so interesting is that the probability $P(p)$ does not change at all with time. This, of course, is a consequence of momentum conservation, but it only holds if there are no forces acting on the particle.

9.4 Velocity and Momentum

So far, I haven't explained the connection between the operator \mathbf{P} and the classical notion of momentum—namely, mass times velocity, or

$$v = p/m. \quad (9.8)$$

What do we mean by the velocity of a quantum mechanical particle? The simplest answer is that we mean the time derivative of the average position $\langle \Psi | \mathbf{X} | \Psi \rangle$:

$$v = \frac{d\langle \Psi | \mathbf{X} | \Psi \rangle}{dt}$$

or, more concretely, in terms of wave functions,

$$v = \frac{d}{dt} \int \psi^*(x, t) x \psi(x, t).$$

Why does $\langle \Psi | \mathbf{X} | \Psi \rangle$ vary with time? Because ψ depends on time, and in fact we know just how. The time dependence of ψ is governed by the time-dependent Schrödinger equation. We could use that fact to work out how $\langle \Psi | \mathbf{X} | \Psi \rangle$ varies with time. I've done it this way—by brute force—and it takes several pages. Fortunately, the abstract methods you learned in earlier lectures make it easier; in fact, we have already done most of the work in Lecture 4. In fact, before we proceed, I recommend that you review Lecture 4, especially Section 4.9, from the beginning to the appearance of Eq. 4.17. To restate Eq. 4.17,

$$\frac{d}{dt}\langle \mathbf{L} \rangle = \frac{i}{\hbar} \langle [\mathbf{H}, \mathbf{L}] \rangle.$$

In words: the time derivative of the expectation value of any observable \mathbf{L} is given by i/\hbar times the expectation value of the commutator of the Hamiltonian with \mathbf{L} . Applying this principle to the velocity v , we find that

$$v = \frac{i}{2m\hbar} \langle [\mathbf{P}^2, \mathbf{X}] \rangle. \quad (9.9)$$

Now, all we have to do is compute the commutator of \mathbf{P}^2 and \mathbf{X} . A couple of simple steps shows that

$$[\mathbf{P}^2, \mathbf{X}] = \mathbf{P} [\mathbf{P}, \mathbf{X}] + [\mathbf{P}, \mathbf{X}] \mathbf{P}. \quad (9.10)$$

This relation can be confirmed by expanding each commutator and spotting some obvious cancellations.

Exercise 9.2: Prove Eq. 9.10 by expanding each side and comparing the results.

The last step uses the standard commutation relation

$$[\mathbf{P}, \mathbf{X}] = -i\hbar.$$

Substituting this into Eq. 9.10 and plugging that result into Eq. 9.9, we find that

$$v = \frac{\langle \mathbf{P} \rangle}{m}$$

or, perhaps more familiarly,

$$\langle \mathbf{P} \rangle = mv. \quad (9.11)$$

We have proved exactly what we set out to prove: the momentum is equal to the mass times the velocity, or, more exactly, the average momentum equals the mass times the velocity.

To get a better idea of what this means, let's suppose the wave function has the form of a packet, or fairly narrow lump. The expectation value of x will be approximately at the center of the lump. What Eq. 9.11 tells us is that the center of the wave packet travels according to the classical rule $p = mv$.

9.5 Quantization

Before moving on to the subject of forces in quantum mechanics, I want to pause and discuss what we have done. We started with a well-known and well-trusted classical system—the free particle—and *quantized* it. We can codify this procedure as follows:

1. Start with a classical system. This means a set of coordinates x and momenta p . In our example, there was only one coordinate and one momentum, but the procedure is easy to generalize. The coordinates and momenta come in pairs, x_i and p_i . The classical system also has a Hamiltonian, which is a function of the x 's and p 's.

2. Replace the classical phase space with a linear vector space. In the position representation, the space of states is represented by a wave function $\psi(x)$ that depends on the coordinates—in general, all of them.
3. Replace the x 's and p 's with operators \mathbf{X}_i and \mathbf{P}_i . Each \mathbf{X}_i acts on the wave function to multiply it by x_i . Each \mathbf{P}_i acts according to the rule

$$\mathbf{P}_i \rightarrow -i\hbar \frac{\partial}{\partial x_i}.$$

4. When these replacements are made, the Hamiltonian becomes an operator that can be used in either the time-dependent or time-independent Schrödinger equation. The time-dependent equation tells us how the wave function changes with time. The time-independent form allows us to find the eigenvectors and eigenvalues of the Hamiltonian.

This procedure of quantization is the means by which the classical equations of a system converted to quantum equations. It has been used over and over, in fields ranging from the motion of particles to quantum electrodynamics; there have even been (not so successful) attempts to quantize Einstein's theory of gravity. As we saw in one simple case, the procedure guarantees that the motion of expectation values is closely related to classical motion.

All of this raises a “chicken and egg” question: Which comes first—classical theory or quantum theory? Should

the logical starting point of physics be classical or quantum mechanical? I think the answer is obvious. Quantum mechanics is the real description of nature. Classical mechanics, while beautiful and elegant, is nevertheless an approximation. Roughly speaking, it holds true when wave functions maintain their shape as packets. Sometimes, we're lucky and the quantum theory of a system can be guessed—and that's all it is, a guess—by starting with a familiar classical system and quantizing it. Sometimes this works. The quantum motion of electrons, deduced from the classical mechanics of particles, is a case in point. Quantum electrodynamics, deduced from Maxwell's equations, is another. But sometimes there is no classical theory to use as a starting point. The spin of a particle has no real classical counterpart. And the quantization of general relativity has largely failed. Quantum theory is probably much more fundamental than classical theory, which generally should be understood as an approximation.

That being said, I will now continue to quantize the motion of particles, but this time incorporating the effects of forces.

9.6 Forces

The world would be a dull place if all particles were free. Forces are what make particles do interesting things, such as assembling themselves into atoms, molecules, chocolate bars, and black holes. The force on any given particle is the sum total of the forces exerted on it by all the other particles in the universe. In practice, we usually assume that we know

what all the other particles are doing and replace their effect by a potential energy function for the particle that we are studying. This much is true in both classical and quantum mechanics.

The potential energy function is denoted by $V(x)$. In classical mechanics it's related to the force on a particle by the equation

$$F(x) = -\frac{\partial V}{\partial x}.$$

If the motion is one-dimensional, the partial derivative can be replaced by an ordinary derivative, but I will leave it as is. If we then combine this equation with Newton's second law, $F = ma$, we get

$$m \frac{d^2 x}{dt^2} = -\frac{\partial V}{\partial x}.$$

In quantum mechanics, we proceed differently; we write a Hamiltonian and solve the Schrödinger equation. Incorporating the potential energy into this program is straightforward. The potential energy $V(x)$ becomes an operator \mathbf{V} that gets added to the Hamiltonian.

What kind of operator is \mathbf{V} ? The answer is easiest to express if we think in the language of wave functions rather than in terms of abstract bras and kets. When the operator \mathbf{V} acts on any wave function $\psi(x)$, it multiplies the wave function by the function $V(x)$.

$$\mathbf{V}|\Psi\rangle \rightarrow V(x)\psi(x).$$

Just as in classical mechanics, once forces are included, the momentum of a particle is not conserved. In fact, Newton's laws of motion can be stated in the form

$$\frac{dp}{dt} = F$$

or

$$\frac{dp}{dt} = -\frac{\partial V}{\partial x}. \quad (9.12)$$

The rules of quantization require us to add $\mathbf{V}(x)$ to the Hamiltonian,⁴

$$\mathbf{H} = \frac{\mathbf{P}^2}{2m} + \mathbf{V}(x), \quad (9.13)$$

and modify the Schrödinger equations in the obvious way:

$$\begin{aligned} i\hbar \frac{\partial \psi}{\partial t} &= \frac{-\hbar^2}{2m} \frac{\partial^2 \psi}{\partial x^2} + V(x)\psi \\ E\psi &= \frac{-\hbar^2}{2m} \frac{\partial^2 \psi}{\partial x^2} + V(x)\psi. \end{aligned} \quad (9.14)$$

What effect does this have? The additional term certainly affects the way ψ changes with time. That of course must be so if the average position of a wave packet is to follow a classical trajectory. To check our reasoning, let's see if it does. First of all, does Eq. 9.11 still hold? It should,

⁴Technically, this is true for free particles as well. However, in the case of free particles we set $V(x)$ equal to 0.

because the connection between momentum and velocity is unaffected by the presence of forces.

Because a new term has been added to \mathbf{H} , there will be a new term in the commutator of \mathbf{X} and \mathbf{H} . Potentially, that could modify the expression for velocity in Eq. 9.9, but it's easy to see that this doesn't happen. The new term involves the commutator of \mathbf{X} with $\mathbf{V}(x)$. But multiplying by x and multiplying by a function of x are operations that commute. In other words,

$$[\mathbf{X}, \mathbf{V}(x)] = 0.$$

Therefore, the connection between velocity and momentum is unaffected by forces in quantum mechanics, as is the case in classical mechanics.

The more interesting question is: Can we understand the quantum version of Newton's law? As stated above, this law can be written as

$$\frac{dp}{dt} = F.$$

Let's calculate the time derivative of the expectation value of \mathbf{P} . Again, the trick is to commute \mathbf{P} with the Hamiltonian:

$$\frac{d}{dt}\langle \mathbf{P} \rangle = \frac{i}{2m\hbar}\langle [\mathbf{P}^2, \mathbf{P}] \rangle + \frac{i}{\hbar}\langle [\mathbf{V}, \mathbf{P}] \rangle. \quad (9.15)$$

The first term is zero because an operator commutes with any function of itself. To compute the second term, we'll use

an equation that we haven't proved yet:

$$[\mathbf{V}(x), \mathbf{P}] = i\hbar \frac{dV(x)}{dx}. \quad (9.16)$$

Plugging Eq. 9.16 into Eq. 9.15, we get

$$\frac{d}{dt} \langle \mathbf{P} \rangle = - \left\langle \frac{dV}{dx} \right\rangle.$$

Now, let's prove Eq. 9.16. Letting the commutator act on a wave function, we can write

$$[\mathbf{V}(x), \mathbf{P}] \psi(x) = V(x) \left(-i\hbar \frac{d}{dx}\right) \psi(x) - \left(-i\hbar \frac{d}{dx}\right) V(x) \psi(x). \quad (9.17)$$

This is easily simplified and results in Eq. 9.16. Thus, we have shown that

$$\frac{d}{dt} \langle \mathbf{P} \rangle = - \left\langle \frac{dV}{dx} \right\rangle, \quad (9.18)$$

which is the quantum analog of Newton's equation for the time rate of change of momentum.

Exercise 9.3: Show that the right-hand side of Eq. 9.17 simplifies to the right-hand side of Eq. 9.16. *Hint:* First expand the second term by taking the derivative of the product. Then look for cancellations.

9.7 Linear Motion and the Classical Limit

You might think we have proved that the expectation value of \mathbf{X} exactly follows the classical trajectory. But what we've actually proved is quite different. This difference exists because the average of a function of x is not the same as the function of the average of x . If Eq. 9.18 had read

$$\frac{d}{dt}\langle\mathbf{P}\rangle = -\frac{dV(\langle x\rangle)}{d\langle x\rangle} \quad [This\ is\ wrong]$$

(and, let me emphasize, *it does not*), then indeed we would say that the average position and momentum satisfy the classical equations. But in reality the classical equations are only approximations, good whenever we can replace the average of dV/dx by the function of the average of x . When is it reasonable to do this? The answer is whenever the $V(x)$ varies slowly compared to the size of the wave packet. If V varies rapidly across the wave packet, the classical approximation will break down. In fact, in that situation a nice, narrow wave packet will get broken up into a badly scattered wave that has no resemblance to the original wave packet. The probability function will also get scattered. Then you'll have no choice but to solve the Schrödinger equation.

Let's look at this point more closely. Mathematically, we've made no assumptions about the shapes of our wave packets. But we have tacitly thought of them as being nicely shaped functions with a single maximum, smoothly trailing off to zero in the positive and negative directions. This con-

dition, though not explicit in our mathematical assumptions, does have a real impact on whether a particle behaves the way classical mechanics would lead us to expect.

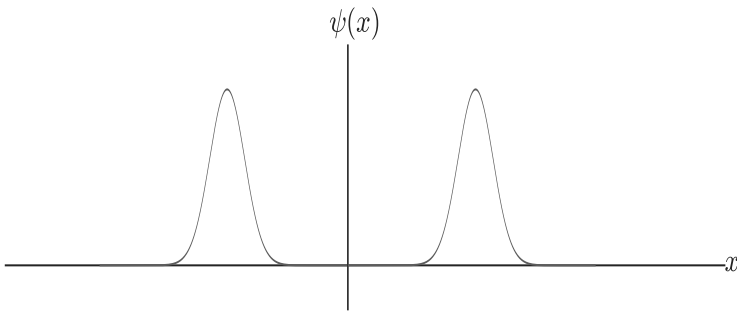


Figure 9.3: Bimodal (Two-Humped) Function, Centered at $x = 0$. Note that $\langle x \rangle = 0$, but $\Delta x > 0$.

To illustrate this point, let's consider a slightly “weird” wave packet. Fig. 9.3 shows a bimodal wave packet (having two maxima), centered at the origin of the x axis. Now, let's consider some function of x , say $F(x)$, where F represents force. The expectation value of $F(x)$ is not the same as the function F of the expectation value of x . In other words,

$$\langle F(x) \rangle \neq F(\langle x \rangle).$$

The right-hand side is a function of the center of the wave packet. It is not the same as the left-hand side, which corre-

sponds to our results from the previous section— $\langle F(x) \rangle$ has the same form as the right-hand side of Eq. 9.18.⁵

Let me give you an example where these two expressions could be extremely different. Suppose that F is equal to x squared:

$$F = x^2.$$

And suppose the wave packet looks like Fig. 9.3. What's the expectation value of x ? It's zero, and so is $F(\langle x \rangle)$, because $F(0) = 0^2 = 0$. On the other hand, what is the expectation value of x^2 ? It's greater than zero. So when a wave packet is not a nice, single bump that is mainly characterized by its center, it's not always true that the time rate of change of the momentum is the force evaluated at the expectation value of x . It's only when the wave function is concentrated over a fairly narrow range that the expectation value of $F(x)$ is the same as $F(\langle x \rangle)$. So we have cheated a little in saying our quantum equation of motion looks classical. That depends on the wave packet being coherent and well localized.

Everything else being equal, when the mass of a particle is large, the wave function tends to be very well concentrated. If there are no very sharp spikes in the potential function $V(x)$, then it will be a good approximation to replace $\langle F(x) \rangle$ with $F(\langle x \rangle)$. When $V(x)$ has spikes, however, the wave packet tends to break up. For example, suppose we have a nice wave packet moving to the right, and it hits a point structure, like an atom, with a potential function sim-

⁵Recall that $-\langle \frac{dV}{dx} \rangle$ represents force in that equation.

ilar to Fig. 9.4. The wave packet will spread out and disintegrate. If, on the other hand, it hits a very smooth potential, then it will go through the smooth potential, moving more or less according to the classical equations of motion. We don't expect quantum mechanics to reproduce classical mechanics in every possible circumstance. We expect it to reproduce classical mechanics in circumstances where it should—where the particles are heavy, the potentials are smooth, and nothing causes the wave function to disintegrate or scatter.⁶

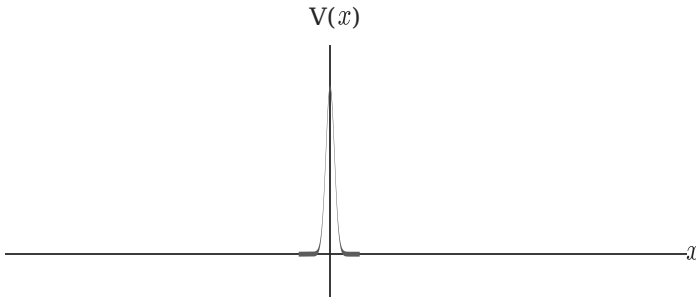


Figure 9.4: Spiky Potential Function. Potential functions with sharp peaks tend to cause wave functions to scatter. The smaller these features are in relation to the wave packet, the more the wave packet will scatter, and the less “classical” it will become.

What physical situations lead to “bad potentials” that break up the wave function? Suppose a potential has fea-

⁶Not as eloquent as Garrison Keillor’s tagline, but true all the same.

tures that have a certain size associated with them. Think of Fig. 9.4 on steroids, with lots of large, closely packed spikes. Suppose we call the size of these features δx , and that δx is significantly smaller than the incoming particle's uncertainty in position:

$$\delta x < \Delta x.$$

If the sharp features of $V(x)$ exist on a scale that is much smaller than the size of the incoming wave packet, the packet will break into a lot of little pieces. Each one will scatter off in a different direction. Roughly speaking, when the features of the potential are shorter than the wavelength of the incoming particle, the wave function will tend to break up.

Let's say you take a bowling ball and ask, "What is Δx ?" We can use the uncertainty principle to gain some intuition about this question. Typically, $\Delta p \times \Delta x$ is bigger than \hbar . But in many reasonable cases it's of order \hbar :

$$\Delta p \Delta x \sim \hbar.$$

Now, p is about as concentrated as it can be, but for an ordinary macroscopic object, the uncertainty relation is pretty much saturated—the left-hand side is roughly *equal* to \hbar . The reasons for this are very complicated, and I won't go into them here. Instead, let's assume this is true and work out the implications. What is Δp ? It's $m\Delta v$, which gives us

$$m\Delta v \Delta x \sim \hbar.$$

Rearranging the symbols, we can then write

$$\Delta v \Delta x \sim \frac{\hbar}{m}$$

or

$$\Delta x \sim \frac{\hbar}{m \Delta v}.$$

Now, if I put a bowling ball on the ground, I know very well that the uncertainty in its velocity is not very big. As the ball gets heavier and heavier, you might expect the uncertainty in velocity to get smaller and smaller. But, in any case, the right-hand side has an m in the denominator, and regardless of Δv , as m gets smaller, Δx will get bigger. And in particular, it will tend to get bigger than the features in the potential.

In the quantum mechanical limit where m is very small and Δx tends to be big, the wave function will move under the influence of a ragged potential, which it sees as being much sharper and more featured than the wave function itself. That's when the wave function breaks up. On the other hand, as m gets very large, Δx gets small. For a large bowling ball, the wave packet might be *very* concentrated. When it moves through a spiky potential, this *tiny* wave function encounters a potential whose features are (comparatively) very broad. Moving through broad smooth features does not break the wave function into pieces. Large masses and smooth potentials characterize the classical limit. A particle with low mass, moving through an abrupt potential, behaves like a quantum mechanical system.

What about electrons? Are they massive enough to behave classically? The answer depends on the interplay between the potential and the mass. For example, if you have two capacitor plates separated by a centimeter, with a smooth electric field between them, then the electron will move across the gap like a nice, coherent, almost classical particle. On the other hand, the potential associated with the nucleus of an atom always has a sharp feature in it. If an electron wave packet hits this potential, it will scatter all over the place.

Before leaving this topic, I'd like to mention *minimum-uncertainty* wave packets. These are wave packets where $\Delta x \Delta p$ is equal to $\hbar/2$ (as opposed to being greater). In other words, in these cases, $\Delta x \Delta p$ is as small as quantum mechanics allows. These wave packets have the form of a Gaussian curve, and they're often called *Gaussian* wave packets. Over time, they spread out and flatten. Such wave packets are not that common, but they do exist. A bowling ball at rest is a good approximation. In Lecture 10, we'll see that the ground state of a harmonic oscillator is a Gaussian wave packet.

9.8 Path Integrals

Classical Hamiltonian mechanics focuses on the step-by-step incremental changes in the state of a system. But there is another way to formulate mechanics—the Principle of Least Action—in which the focus is on entire histories. For a particle, this means looking at the full trajectory of the particle from some initial time to some final time. The content of the two approaches is the same, but the emphasis is differ-

ent. Hamiltonian mechanics zeros in on some instant and tells you how the system changes between that instant and the next. The least action principle steps back and takes a global look. One can imagine nature sampling all possible trajectories and picking the one that minimizes the action between a pair of fixed initial and final points.⁷

Quantum mechanics also has a Hamiltonian description that concentrates on incremental changes. It's called the time-dependent Schrödinger equation, and it's very general. As far as we know, it can be used to describe all physical systems. Still, it seems fair to ask, as Richard Feynman did almost seventy years ago, whether there is a way to look at quantum mechanics that pictures whole histories. In other words, is there a formulation that parallels the Principle of Least Action? I will not explain Feynman's path integral description in detail in this lecture, but just to whet your appetite I'll give you a hint of how it works.

First, let me very briefly remind you of the classical least action principle as I explained it in *Volume I*. Suppose that a classical particle starts at position x_1 at time t_1 and arrives at position x_2 at time t_2 (Fig. 9.5). The question is: What is the trajectory that it took between t_1 and t_2 ?

According to the least action principle, the actual trajectory is the one of minimum action. *Action* is of course a technical term, and it stands for the integral of the Lagrangian between the end points of the trajectory. For simple systems,

⁷Strictly speaking, the principle should be called the Principle of Stationary Action. Actual trajectories are stationary points of the action and not always minima. For our purposes, this fine point is not important.

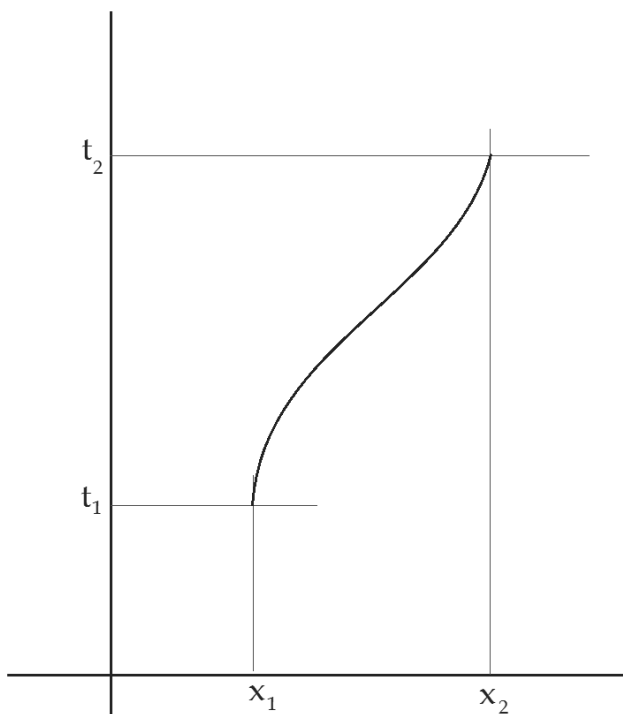


Figure 9.5: Classical Trajectory. This shows one path a particle may take when moving from point 1 (x_1, t_1) to point 2 (x_2, t_2) . To keep things simple, the \dot{x} axis, representing the particle's velocity in the x direction, is not shown.

the Lagrangian is the kinetic energy minus the potential energy. Thus, for a particle that moves in one dimension, the action is

$$A = \int_{t_1}^{t_2} L(x, \dot{x}) dt \quad (9.19)$$

or

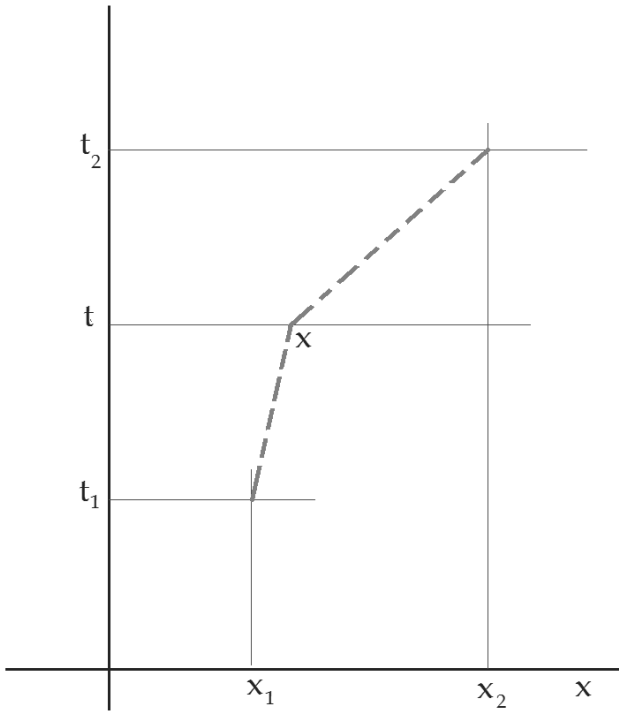


Figure 9.6: First Step Toward Quantizing the Trajectory. Break the particle's path into two equal parts (equal in time, that is). The particle has the same starting and ending points, but now its trajectory passes through the intermediate point x .

$$A = \int_{t_1}^{t_2} \left(\frac{m\dot{x}^2}{2} - V(x) \right) dt.$$

The idea is to try out all possible trajectories connecting the two end points, and calculate A for each one of them. The

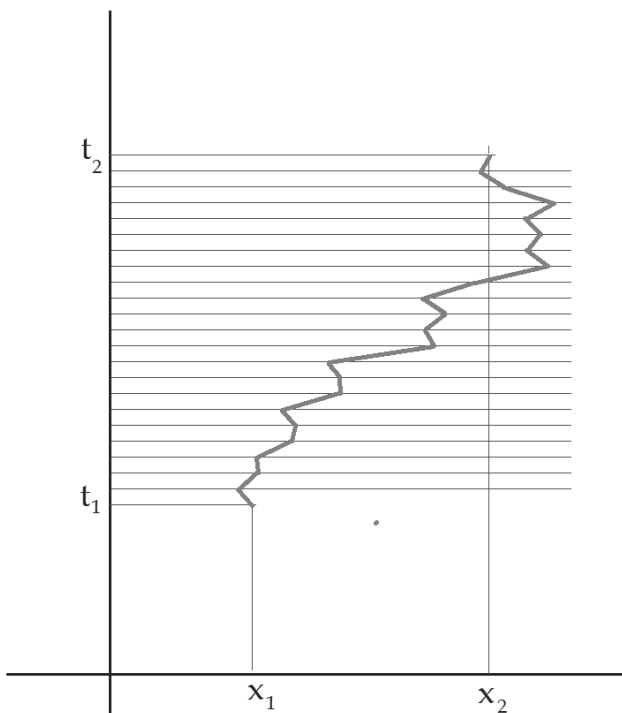


Figure 9.7: Further Steps Toward Constructing the Path Integral. Keeping the same starting and ending points, break the path up into a large number of equally sized segments.

winner is the one that has the least action.^{8,9}

Now, let's turn to quantum mechanics. The idea of a well-defined trajectory between two points makes no sense in quantum mechanics because of the uncertainty principle.

⁸That's how it works conceptually, anyway. In practice, the Euler-Lagrange equations provide a shortcut, as explained in *Volume I*.

⁹To keep our diagrams simple, we don't display an \dot{x} axis even though the Lagrangian clearly depends on \dot{x} .

However, a question that we *can* ask is: Given that a particle starts out at (x_1, t_1) , what is the probability that it will show up at (x_2, t_2) if an observation of its position is made?

As always in quantum mechanics, the probability is the square of the absolute value of a complex amplitude. The global version of quantum mechanics asks:

Given that a particle starts out at (x_1, t_1) , what is the amplitude that it will show up at (x_2, t_2) ?

Let's call that amplitude $C(x_1, t_1; x_2, t_2)$ or, more simply, just $C_{1,2}$. The initial state of the particle is $|\Psi(t_1)\rangle = |x_1\rangle$. Over the time interval between t_1 and t_2 , the state evolves to

$$|\Psi(t_2)\rangle = e^{-iH(t_2-t_1)}|x_1\rangle. \quad (9.20)$$

The amplitude to detect the particle at $|x_2\rangle$ is just the inner product of $|\Psi(t_2)\rangle$ with $|x_2\rangle$. Its value is

$$C_{1,2} = \langle x_2 | e^{-iH(t_2-t_1)} | x_1 \rangle. \quad (9.21)$$

In other words, the amplitude to go from x_1 to x_2 over the time interval $t_2 - t_1$ is constructed by sandwiching $e^{-iH(t_2-t_1)}$ between the initial and final positions. To simplify the formula, let's define $t_2 - t_1$ to be t . Then the amplitude is

$$C_{1,2} = \langle x_2 | e^{-iHt} | x_1 \rangle. \quad (9.22)$$

Now, let's break the time interval t into two smaller intervals of size $t/2$ (see Fig. 9.6). The operator e^{-iHt} can be written

as the product of two operators:

$$e^{-iHt} = e^{-iHt/2}e^{-iHt/2}. \quad (9.23)$$

By inserting the identity operator in the form

$$I = \int dx |x\rangle\langle x|, \quad (9.24)$$

we can rewrite the amplitude as

$$C_{1,2} = \int dx \langle x_2 | e^{-iHt/2} | x \rangle \langle x | e^{-iHt/2} | x_1 \rangle. \quad (9.25)$$

This form of the equation looks more complicated, but has a very interesting interpretation. Let me put it in words. The amplitude to get from x_1 to x_2 over time interval t is an integral over an intermediate position x . The integrand is the amplitude to go from x_1 to x over the time interval $t/2$ *multiplied by* the amplitude to go from x to x_2 over another time interval $t/2$.

Fig. 9.6 shows the same idea in visual terms. Classically, to go from x_1 to x_2 , the particle must pass through an intermediate point x . But in quantum mechanics the amplitude to go from x_1 to x_2 is an integral over all possible intermediate points.

We can carry this idea further and divide the time interval into a great many tiny intervals, as illustrated in Fig. 9.7. I won't write out the complicated formulas, but the idea should be clear. For each tiny time interval, say of size e , we include a factor

$$e^{-i\epsilon H}.$$

Then, between each pair of factors, we insert the identity so that the amplitude $C_{1,2}$ becomes a multiple integral over all the intermediate locations. The integrand is built from products of expressions with the form

$$\langle x_i | e^{-i\epsilon H} | x_{i+1} \rangle.$$

If we define $U(\epsilon)$ as

$$U(\epsilon) = e^{-i\epsilon H},$$

then we can write the entire product as

$$\langle x_2 | U^N | x_1 \rangle$$

or

$$\langle x_2 | UUUU \dots | x_1 \rangle.$$

In this equation, U appears N times as a factor, where N is the number of epsilon steps. We can then insert identity operators between the U 's.

Such an expression can be called the amplitude for the given path. But the particle does not travel along a particular path. Instead, in the limit of a large number of infinitesimal time intervals, *the amplitude is an integral over*

all possible paths between the end points. The elegant fact that Feynman discovered is that the amplitude for each path bears a simple relation to a familiar expression from classical mechanics—the action for that path. The exact expression for each path is

$$e^{iA/\hbar},$$

where A is the action for the individual path.

Feynman's formulation can be summarized by a single equation:

$$C_{1,2} = \int_{paths} e^{iA/\hbar}. \quad (9.26)$$

The path integral formulation is not merely an elegant mathematical trick; it has real power. In fact, it can be used to derive both Schrödinger equations, and all the commutation relations of quantum mechanics. But it really comes into its own in the context of quantum field theory, where it is the principal tool for formulating the laws of elementary particle physics.

Lecture 10

The Harmonic Oscillator

Art: I think I see it, Lenny. The whole picture is slowly coming into focus. Minus One, General Uncertainty, entangled pairs, the Hamiltonian—even the degenerates. What’s next?

Lenny: Oscillations, Art. Vibrations. You’re a fiddler—play us a last tune tonight. Something with good vibes.

Of all the ingredients that go into building a quantum description of the world, two stand out as especially fundamental. The spin, or qubit, of course is one of them. In classical logic, everything can be built out of yes-no questions. Similarly, in quantum mechanics, every logical question boils down to a question about qubits. We spent a lot of time in earlier lectures learning about qubits. In this lecture, we’ll learn about the second basic ingredient of quantum mechanics—the harmonic oscillator.

The harmonic oscillator isn’t a particular object like a hydrogen atom or a quark. It’s really a mathematical framework for understanding a huge number of phenomena. This

concept of the harmonic oscillator also exists in classical physics, but it really comes to the fore in quantum theory.

One example of a harmonic oscillator is a particle moving under a linear restoring force; for example, the iconic weight on the end of a spring. An idealized spring satisfies Hooke's law: the force on the displaced mass is proportional to the distance it has been displaced. We call the force a *restoring force* because it pulls the mass back toward the equilibrium position.

Another example is a marble rolling back and forth at the bottom of a bowl, with no energy being lost to friction. What characterizes these systems is a potential energy function that looks like a parabola:

$$V(x) = \frac{k}{2}x^2. \quad (10.1)$$

The constant k is called the *spring constant*. If we recall that the force on an object is minus the gradient of V , we find that the force on the object is

$$F = -kx. \quad (10.2)$$

The negative sign tells us that the force acts opposite to the displacement and pulls the mass back toward the origin.

Why are harmonic oscillators so prevalent in physics? Because almost any smooth function looks like a parabola close to a minimum of the function. Indeed, many kinds of systems are characterized by an energy function that can be approximated by a quadratic function of some variable representing a displacement from equilibrium. When disturbed,

these systems will all oscillate about the equilibrium point. Here are some other examples:

- An atom situated in a crystal lattice. If the atom is displaced slightly from its equilibrium position, it gets pushed back with an approximately linear restoring force. This motion is three-dimensional and really consists of three independent oscillations.
- The electric current in a circuit of low resistance often oscillates with a characteristic frequency. The mathematics of circuits is identical to the mathematics of masses attached to springs.
- Waves. If the surface of a pond is disturbed, it sends out waves. Someone watching at a particular location will see the surface oscillate as the wave passes by. This motion can be described as simple harmonic motion. The same goes for sound waves.
- Electromagnetic waves. Just like any other wave, a light wave or a radio wave oscillates when it passes you. The same mathematics that describes the oscillating particle also applies to electromagnetic waves.

The list goes on and on but the math is always the same. Just to have an example in mind, let's picture the oscillator as a weight hanging from a spring. Needless to say, we hardly need quantum mechanics to describe an ordinary weight and spring, so let's imagine a very tiny version of this same system and then quantize it.

10.1 The Classical Description

Let's use y to denote the height of the hanging weight. We'll choose the origin so that the weight is at $y = 0$ when it's in equilibrium—that is when the weight is hanging at rest. To study this system classically, we can use the Lagrangian method that we learned about in *Volume I*. The kinetic and potential energies are $\frac{1}{2}m\dot{y}^2$ and $\frac{1}{2}ky^2$ respectively.

As you recall, the Lagrangian is the kinetic energy *minus* the potential energy:

$$L = \frac{1}{2}m\dot{y}^2 - \frac{1}{2}ky^2.$$

First, we'll put the Lagrangian into a certain standard form by changing from y to another variable that we will call x . This coordinate is not something new. It still represents the displacement of the mass. By switching from y to x , we're just making a convenient change of units. Let's define the new variable as

$$x = \sqrt{m} y.$$

In terms of x , the Lagrangian becomes

$$L = \frac{1}{2}\dot{x}^2 - \frac{1}{2}\omega^2 x^2. \quad (10.3)$$

The constant ω is defined as $\omega = \sqrt{\frac{k}{m}}$ and happens to be the frequency of the oscillator.

By making this change of variables, we can describe every oscillator in exactly the same form. In this form, oscillators are distinguished from each other only by their frequency ω .

Now, let's use Lagrange's equations to work out the equations of motion. For this one-dimensional system, there is only one Lagrange equation, namely

$$\frac{\partial L}{\partial x} = \frac{d}{dt} \frac{\partial L}{\partial \dot{x}}. \quad (10.4)$$

Carrying out these operations on Eq. 10.3, we find that

$$\frac{\partial L}{\partial \dot{x}} = \dot{x}. \quad (10.5)$$

This is called the canonical momentum conjugate to x . Differentiating with respect to time gives

$$\frac{d}{dt} \frac{\partial L}{\partial \dot{x}} = \ddot{x}, \quad (10.6)$$

and now we have the right-hand side of Eq. 10.4. Turning to the left-hand side, we find that

$$\frac{\partial L}{\partial x} = -\omega^2 x. \quad (10.7)$$

Setting the left and right sides (Eqs. 10.7 and 10.6) of the Lagrange equation equal to each other, we get

$$-\omega^2 x = \ddot{x}. \quad (10.8)$$

This equation is, of course, equivalent to $F = ma$. Why is there a minus sign? Because the force is a restoring force—its direction is opposite to the direction of the displacement. By now you have seen this type of equation enough to know that the solution contains sines and cosines. The general solution is

$$x = A \cos(\omega t) + B \sin(\omega t), \quad (10.9)$$

which shows us that ω is indeed the frequency of the oscillator. When we differentiate twice, we pull out a factor of ω^2 .

Exercise 10.1: Find the second time derivative of x in Eq. 10.9, and thereby show that it solves Eq. 10.8.

10.2 The Quantum Mechanical Description

Now, let's return to our microscopic version of the weight-and-spring system—let's say no bigger than a single molecule. At first, this seems ridiculous. How could we ever build a spring that small? But in fact nature provides all sorts of microscopic springs. Many molecules consist of two atoms—for example, a heavy atom and a light one. There are forces holding the molecule in equilibrium with the atoms separated by a certain distance. When the light atom is displaced, it will be attracted back to the equilibrium location. The

molecule is a miniature version of the weight-and-spring system, but is so small that we have to use quantum mechanics to understand it.

Having worked out the classical Lagrangian, let's try to build a quantum mechanical description of our system. The first thing we need is a space of states. As we've seen, the state of a particle moving on a line is represented by a wave function $\psi(x)$. There are many possible system states, and each one is represented by a different wave function. A function $\psi(x)$ is defined in such a way that $\psi^*(x)\psi(x)$ is the probability density (the probability per unit interval) to find the particle at position x :

$$\psi^*(x)\psi(x) = P(x).$$

In this equation, $P(x)$ represents the probability density. We now have a sort of kinematics—a specification of what the system states are.

Can $\psi(x)$ be any function at all? Aside from the requirement that it must be continuous and differentiable, the only extra condition is that the total probability of finding the particle at any position must be 1:

$$\int_{-\infty}^{+\infty} \psi^*(x)\psi(x)dx = 1. \quad (10.10)$$

This would not seem to be much of a restriction. Whatever the right-hand side of this equation is, we could always multiply ψ by some constant to make the integral equal to 1—unless the integral is either zero or infinity. Since $\psi^*(x)\psi(x)$

is positive, we don't have to worry about zero, but infinity is a different matter altogether; there are lots of functions that would make the integral in Eq. 10.10 blow up. The conditions for a sensible wave function thus include the requirement that ψ falls to zero fast enough that the integral converges. Functions that meet this condition are called *normalizable*.

There are two questions we might ask about our harmonic oscillator:

- How does the state-vector change as a function of time? To answer this question, we need to know the Hamiltonian.
- What are the oscillator's possible energies? These are also determined by the Hamiltonian.

So to know anything useful we need the Hamiltonian. Fortunately, we can derive it from the Lagrangian, and I'll remind you how in a moment. But first recall that the canonical momentum conjugate to x is defined as $\partial L/\partial \dot{x}$.¹ Combining this with Eq. 10.5, we get

$$p = \frac{\partial L}{\partial \dot{x}} = \dot{x}.$$

Using the straightforward definition from classical mechanics, we find that the Hamiltonian for the harmonic oscillator is

¹This idea is explained in *Volume I*.

$$H = p\dot{x} - \mathcal{L},$$

where p is the canonical momentum conjugate to x , and \mathcal{L} represents the Lagrangian.² We could work directly from this definition, but instead we'll take a shortcut. Because the Lagrangian is the kinetic energy *minus* the potential energy, the Hamiltonian is the kinetic energy *plus* the potential energy—in other words, the total energy. The Hamiltonian for the oscillator can therefore be written

$$H = \frac{1}{2}\dot{x}^2 + \frac{1}{2}\omega^2 x^2.$$

So far, so good, but we're not quite finished. We've expressed kinetic energy in terms of velocity; in quantum mechanics, however, we need to represent our observables as *operators*, and we don't have a velocity operator. To take care of this, we'll have to recast things in terms of position and canonical *momentum*, which does have a standard operator form. Rewriting the Hamiltonian in terms of canonical momentum is easy because

$$p = \frac{\partial L}{\partial \dot{x}} = \dot{x},$$

which allows us to write

$$H = \frac{1}{2}p^2 + \frac{1}{2}\omega^2 x^2. \quad (10.11)$$

²We don't need to use a summation sign because there's only one degree of freedom.

That's the classical Hamiltonian. We can now turn it into a quantum mechanical equation by reinterpreting x and p as operators, defined by their action on $\psi(x)$. As we've done before, we'll use the boldface symbols, \mathbf{X} and \mathbf{P} , to distinguish our quantum operators from their classical counterparts, x and p . From previous lectures, we know exactly how these operators work. \mathbf{X} just multiplies the wave function by the position variable:

$$\mathbf{X}|\psi(x)\rangle \quad \Longrightarrow \quad x\psi(x).$$

And \mathbf{P} takes the same form it does for other one-dimensional problems:

$$\mathbf{P}|\psi(x)\rangle \quad \Longrightarrow \quad -i\hbar \frac{d}{dx}\psi(x).$$

Now, we can figure out the action of the Hamiltonian on a wave function by letting \mathbf{P} act twice on the wave function. This is the same procedure we followed in Lecture 9. In other words,

$$\mathbf{H}|\psi(x)\rangle \quad \Longrightarrow \quad \frac{1}{2}\left(-i\hbar \frac{\partial}{\partial x}\left(-i\hbar \frac{\partial \psi(x)}{\partial x}\right)\right) + \frac{1}{2}\omega^2 x^2 \psi(x),$$

or

$$\mathbf{H}|\psi(x)\rangle \quad \Longrightarrow \quad -\frac{\hbar^2}{2} \frac{\partial^2 \psi(x)}{\partial x^2} + \frac{1}{2}\omega^2 x^2 \psi(x). \quad (10.12)$$

We're using partial derivatives because in general ψ also depends on another variable, *time*. Time is not an operator and

does not have the same status as x , but the state-vector does change with time, and we therefore treat time as a parameter. The partial derivative indicates that we're describing the system "at a fixed time."

10.3 The Schrödinger Equation

Eq. 10.12 shows how the Hamiltonian operates on ψ . Now, let's put it to work. As we said in the previous section, one of its jobs is to tell you how the state-vector changes with time. So let's write out the time-dependent Schrödinger equation:

$$i\frac{\partial\psi}{\partial t} = \frac{1}{\hbar}\mathbf{H}\psi.$$

Substituting for \mathbf{H} using 10.12, we get

$$i\frac{\partial\psi}{\partial t} = -\frac{\hbar}{2}\frac{\partial^2\psi}{\partial x^2} + \frac{1}{2\hbar}\omega^2x^2\psi. \quad (10.13)$$

This equation says that if you know ψ (both the real and imaginary parts) at some particular time, you can predict what it will be at a future time. Notice that the equation is complex—it contains i as a factor. This means that even if ψ starts out being real-valued at time $t = 0$, it will very shortly develop an imaginary part. Any solution ψ must therefore be a complex function of x and t .

You can solve this equation in a number of ways. For example, you can solve it numerically on a computer. Start with a known value of $\psi(x)$ and update it slightly by calculating the derivative. Once you have the derivative, calculate how $\psi(x)$ changes in a small increment of time. Then,

add this incremental change to $\psi(x)$ and keep doing it over and over. It turns out that $\psi(x)$ will do some interesting things—it will move around somehow. In fact, under certain circumstances, it will form a wave packet that moves around very much like a harmonic oscillator.

10.4 Energy Levels

The other thing you can do with the Hamiltonian is calculate the energy levels of the oscillator, by finding the energy eigenvectors and eigenvalues. As we learned in Lecture 4, once you know these eigenvectors and eigenvalues, you can figure out the time dependence without solving any differential equations. That's because you already know the time dependence of each energy eigenvector. You may want to review the Schrödinger's Ket recipe we gave in Section 4.13.

For now, let's concentrate on finding the energy eigenvectors themselves, using the time-independent Schrödinger equation:

$$\mathbf{H}|\psi_E\rangle = E|\psi_E\rangle.$$

The subscript E indicates that ψ_E is the eigenvector for a particular eigenvalue E . This equation defines two things: the wave functions $\psi_E(x)$ and the energy levels E . Let's make things less abstract by expanding \mathbf{H} using Eq. 10.12:

$$-\frac{\hbar^2}{2} \frac{\partial^2 \psi_E(x)}{\partial x^2} + \frac{1}{2} \omega^2 x^2 \psi_E(x) = E \psi_E(x). \quad (10.14)$$

To solve this equation, we must:

- Find the allowable values of E that permit a mathematical solution.
- Find the eigenvectors and possible eigenvalues of the energy.

This is a little trickier than you might think. There turns out to be a solution to the equation for every value of E , including all the complex numbers, but most solutions are physically absurd. If we just start at some point and solve the Schrödinger equation by making little incremental steps, we will almost always find that $\psi(x)$ grows or “blows up” as x becomes large. In other words, we may be able to find solutions to the equation, but only very rarely will we find a normalizable solution.

In fact, for most values of E , including all the complex numbers, the solutions of Eq. 10.14 grow exponentially as x approaches ∞ , $-\infty$, or both. This type of solution makes no physical sense; it tells us that there is an overwhelming probability that the oscillator coordinate is infinitely far away. Clearly, we want to impose some condition that gets rid of such solutions. So let’s impose one:

Physical solutions of the Schrödinger equation must be normalizable.

This is a very powerful constraint. In fact, for almost all values of E , there are no normalizable solutions. But for certain very special values of E such solutions do exist, and we will find them.

10.5 The Ground State

What is the lowest possible energy level for a harmonic oscillator? In classical physics, the energy can never be negative because the Hamiltonian has an x^2 term and a p^2 term; to minimize energy, we just set p and x equal to zero. But in quantum mechanics, that's asking too much. The uncertainty principle says that you *can't* set both x and p equal to zero. The best you can do is find a compromise state in which x and p are not too spread out. Because you have to compromise, the lowest possible energy will *not* be zero. Neither p^2 nor x^2 will be zero. Because the operators \mathbf{X}^2 and \mathbf{P}^2 can have only positive eigenvalues, the harmonic oscillator has no negative energy levels, and in fact, it has no state with zero energy either.

If all the energy levels of a system must be positive, there must be a lowest allowable energy and a wave function to go with it. This lowest energy level is called the *ground state* and is denoted by $\psi_0(x)$. Keep in mind that the subscript 0 does not mean that the energy is zero; it means that it is the lowest allowable energy.

There is a very useful mathematical theorem that helps identify the ground state. We won't prove it here, but it is very simple to state:

The ground-state wave function for any potential has no zeros and it's the only energy eigenstate that has no nodes.

So all we have to do to find the ground state of our harmonic oscillator is to find a nodeless solution for some value

of E . It doesn't matter how we find it—we can use mathematical tricks, make guesses, or just ask the professor. Let's use the latter method. (I'll play the role of the professor.)

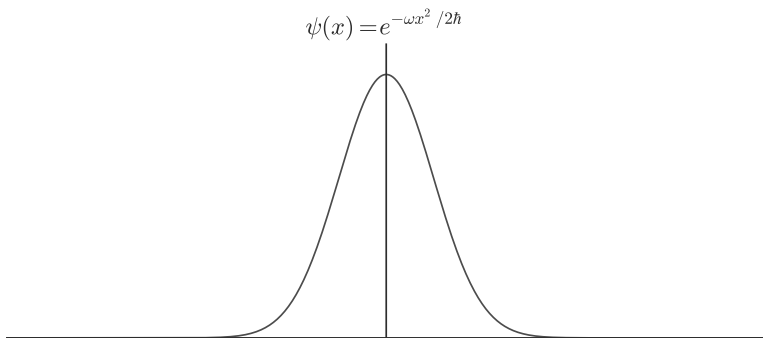


Figure 10.1: Harmonic Oscillator Ground State

Here is a function that works:

$$\psi(x) = e^{-\frac{\omega}{2\hbar}x^2}. \quad (10.15)$$

This function is shown schematically in Fig. 10.1. As you can see, it's concentrated near the origin, where we expect the lowest energy state to be concentrated. It goes to zero very quickly as it moves away from the origin, so the integral of the probability density is finite. And, importantly, it has no nodes. So it has a chance of being our ground state.

Let's see if we can figure out what the Hamiltonian does to this function. The first term of the Hamiltonian (the left side of Eq. 10.14) tells us to apply the operator

$$-\frac{\hbar^2}{2} \frac{\partial}{\partial x^2}$$

to $\psi(x)$. Let's calculate that term, one derivative at a time. The first step is

$$\frac{\partial \psi(x)}{\partial x} = -\frac{\omega}{2\hbar} (2x) e^{-\frac{\omega}{2\hbar} x^2},$$

which simplifies to

$$\frac{\partial \psi(x)}{\partial x} = -\frac{\omega}{\hbar} x e^{-\frac{\omega}{2\hbar} x^2}.$$

When we take the second derivative, there will be two terms because of the product rule:

$$\frac{\partial^2 \psi(x)}{\partial x^2} = -\frac{\omega}{\hbar} e^{-\frac{\omega}{2\hbar} x^2} + \frac{\omega^2}{\hbar^2} x^2 e^{-\frac{\omega}{2\hbar} x^2}.$$

Let's plug this result back into Eq. 10.14, and at the same time replace ψ on the right side with our guess, $e^{-\frac{\omega}{2\hbar} x^2}$:

$$\frac{\hbar}{2} \omega e^{-\frac{\omega}{2\hbar} x^2} - \frac{1}{2} \omega^2 x^2 e^{-\frac{\omega}{2\hbar} x^2} + \frac{1}{2} \omega^2 x^2 e^{-\frac{\omega}{2\hbar} x^2} = E e^{-\frac{\omega}{2\hbar} x^2}.$$

After canceling the terms proportional to $x^2 e^{-\frac{\omega}{2\hbar} x^2}$, we discover the remarkable fact that solving the Schrödinger equation just reduces to solving

$$\frac{\hbar}{2} \omega e^{-\frac{\omega}{2\hbar} x^2} = E e^{-\frac{\omega}{2\hbar} x^2}.$$

As you can see, the only way we can solve this equation is to set the energy E equal to $\frac{\omega\hbar}{2}$. In other words, we've found not only the wave function but also the value of the ground-state energy. Calling the ground-state energy E_0 , we can write

$$E_0 = \frac{\omega\hbar}{2}. \quad (10.16)$$

The ground-state wave function, meanwhile, is just the Gaussian function the professor gave us:

$$\psi_0(x) = e^{-\frac{\omega}{2\hbar}x^2}.$$

He's a clever fellow, that professor.

10.6 Creation and Annihilation Operators

Over the course of these lectures, we have seen two ways of thinking about quantum mechanics. They go all the way back to Heisenberg and Schrödinger. Heisenberg liked algebra, matrices, and, had he known what to call them, linear operators. Schrödinger, by contrast, thought in terms of wave functions and wave equations, the Schrödinger equation being one famous example. Of course, the two ways of thinking are not contradictory; functions form a vector space and derivatives are operators.

So far, in our study of the harmonic oscillator we have focused on functions and differential equations. But the more powerful tool in many cases—particularly for the harmonic

oscillator—is the operator method. It reduces the entire study of wave functions and wave equations to a very small number of algebraic tricks, which almost always involve the commutation relations. In fact, whenever you see a pair of operators, my advice is to figure out their commutator. If the commutator is a new operator that you haven't seen before, find its commutator with the original pair. That's when the fun happens.

Obviously, this advice can lead to an unending chain of boring computations. But once in a while you may get lucky and find a set of operators that *close under commutation*. Whenever that happens, you're in business; as we will see, operator methods have tremendous power.

Now, let's apply this approach to our harmonic oscillator. We begin with the Hamiltonian expressed in terms of the operators \mathbf{P} and \mathbf{X} :

$$\mathbf{H} = \frac{\mathbf{P}^2 + \omega^2 \mathbf{X}^2}{2}. \quad (10.17)$$

To figure out the rest of the energy levels, we'll use some tricks. The idea is to cleverly use the properties of \mathbf{X} and \mathbf{P} (in particular, the commutation relation $[\mathbf{X}, \mathbf{P}] = i\hbar$) to construct two new operators, called *creation* and *annihilation* operators. When a creation operator acts on an energy eigenvector (or eigenfunction), it produces a new eigenvector that has the next higher energy level. An annihilation operator does just the opposite: it produces an eigenvector whose energy is one level lower than the energy of the eigenvector it started with. So, roughly speaking, the thing that they create and annihilate is energy. They're also called *raising*

and *lowering* operators. But remember: operators act on state vectors, not on systems. To see how these operators work, let's rewrite the Hamiltonian in the form

$$\mathbf{H} = \frac{1}{2}(\mathbf{P}^2 + \omega^2\mathbf{X}^2). \quad (10.18)$$

This is a classical as well as a quantum mechanical Hamiltonian, and it would be just as correct to use the lowercase symbols p and x . However, we're using the boldface \mathbf{P} and \mathbf{X} because we plan to focus on the quantum mechanical Hamiltonian.

Let's start by doing a manipulation that is correct for classical physics but will require some modification for quantum mechanics. In the parentheses above, we have a sum of squares. Using the formula

$$a^2 + b^2 = (a + ib)(a - ib),$$

it seems that we can rewrite the Hamiltonian as

$$\mathbf{H} \text{ " = " } \frac{1}{2}(\mathbf{P} + i\omega\mathbf{X})(\mathbf{P} - i\omega\mathbf{X}), \quad (10.19)$$

and that's almost correct. Why *almost*? Because quantum mechanically, \mathbf{P} and \mathbf{X} do not commute, and we need to be careful about the order of operations. Let's expand our factored expression and see how it might differ from the original Hamiltonian in Eq. 10.18. Keeping careful track of the order of factors, we can expand the expression as follows:

$$\begin{aligned}
\frac{1}{2}(\mathbf{P} + i\omega\mathbf{X})(\mathbf{P} - i\omega\mathbf{X}) &= \frac{1}{2}(\mathbf{P}^2 + i\omega\mathbf{XP} - i\omega\mathbf{PX} - i^2\omega^2\mathbf{X}^2) \\
&= \frac{1}{2}(\mathbf{P}^2 + i\omega(\mathbf{XP} - \mathbf{PX}) - i^2\omega^2\mathbf{X}^2) \\
&= \frac{1}{2}(\mathbf{P}^2 + i\omega(\mathbf{XP} - \mathbf{PX}) + \omega^2\mathbf{X}^2) \\
&= \frac{1}{2}(\mathbf{P}^2 + \omega^2\mathbf{X}^2) + \frac{1}{2}i\omega(\mathbf{XP} - \mathbf{PX}).
\end{aligned}$$

Look at the right-hand set of parentheses in the final line. We have seen that expression before—it's the commutator of \mathbf{X} and \mathbf{P} . In fact, we already know its value:

$$(\mathbf{XP} - \mathbf{PX}) = [\mathbf{X}, \mathbf{P}] = i\hbar.$$

Thus, the expression for our factored Hamiltonian becomes

$$\frac{1}{2}(\mathbf{P}^2 + \omega^2\mathbf{X}^2) + \frac{1}{2}i\omega i\hbar$$

or

$$\frac{1}{2}(\mathbf{P}^2 + \omega^2\mathbf{X}^2) - \frac{1}{2}\omega\hbar.$$

In other words, the factored expression we started out with in Eq. 10.19 is actually smaller than the Hamiltonian by $\frac{\omega\hbar}{2}$. To recover the actual Hamiltonian, we need to add the $\frac{\omega\hbar}{2}$ back in:

$$\mathbf{H} = \frac{1}{2}(\mathbf{P} + i\omega\mathbf{X})(\mathbf{P} - i\omega\mathbf{X}) + \frac{\omega\hbar}{2}.$$

Rewriting the Hamiltonian this way and that way may seem like an exercise in futility, but trust me, it's not. First of all, the last term is just an additive constant that adds the numerical value $\frac{\omega\hbar}{2}$ to every energy eigenvalue. We can ignore it for now. Later, after we've solved the rest of the problem, we can add it back in. The guts of the problem are found in the expression $(\mathbf{P} + i\omega\mathbf{X})(\mathbf{P} - i\omega\mathbf{X})$. It turns out that these two factors, $(\mathbf{P} + i\omega\mathbf{X})$ and $(\mathbf{P} - i\omega\mathbf{X})$, have some very remarkable properties. In fact, they are the raising and lowering operators (or creation and annihilation operators) that I told you about earlier. For now, these are just names, but as we go along we'll see that the names were well chosen. The obvious definitions would be

$$\mathbf{a}^- = (\mathbf{P} - i\omega\mathbf{X})$$

for the lowering operator, and

$$\mathbf{a}^+ = (\mathbf{P} + i\omega\mathbf{X})$$

for the raising operator. But history sometimes preempts the obvious. Historically, the raising and lowering operators have been defined with an extra factor in front of them. Here are the official definitions:

$$\mathbf{a}^- = \frac{i}{\sqrt{2\omega\hbar}}(\mathbf{P} - i\omega\mathbf{X}), \quad (10.20)$$

$$\mathbf{a}^+ = \frac{-i}{\sqrt{2\omega\hbar}}(\mathbf{P} + i\omega\mathbf{X}), \quad (10.21)$$

If we use these definitions, the Hamiltonian starts to look very simple:

$$\mathbf{H} = \omega\hbar(\mathbf{a}^+\mathbf{a}^- + 1/2). \quad (10.22)$$

There are only two properties of \mathbf{a}^+ and \mathbf{a}^- that we need to know. The first is that they are Hermitian conjugates of each other. That follows from their definitions. The other property is what really gives them juice. The commutator of \mathbf{a}^+ and \mathbf{a}^- is

$$[\mathbf{a}^-, \mathbf{a}^+] = 1.$$

This is easy to prove. First, we use the definitions to write

$$[\mathbf{a}^-, \mathbf{a}^+] = \frac{1}{2\omega\hbar}[(\mathbf{P} - i\omega\mathbf{X}), (\mathbf{P} + i\omega\mathbf{X})]$$

The next step is to use the commutation relations $[\mathbf{X}, \mathbf{X}] = 0$, $[\mathbf{P}, \mathbf{P}] = 0$, and $[\mathbf{X}, \mathbf{P}] = i\hbar$. Apply these to the above equation, and you will quickly find that $[\mathbf{a}^-, \mathbf{a}^+] = 1$.

We can make the Hamiltonian in Eq. 10.22 even simpler by defining a new operator,

$$\mathbf{N} = \mathbf{a}^+\mathbf{a}^-,$$

called the *number operator*. Once again, this is just a name, but as we'll see, it's a very good name. Stated in terms of the number operator, the Hamiltonian becomes

$$\mathbf{H} = \omega\hbar(\mathbf{N} + 1/2). \quad (10.23)$$

So far, all we've done is define some symbols, \mathbf{a}^+ , \mathbf{a}^- , and \mathbf{N} , that make the Hamiltonian look deceptively simple; it's not clear that we are actually any closer to figuring out the energy eigenvalues. To proceed further, let's recall my earlier advice: whenever you see two operators, commute them. In this case, we already know one commutator:

$$[\mathbf{a}^-, \mathbf{a}^+] = 1. \quad (10.24)$$

Next, let's find the commutator of the raising and lowering operators with the number operator \mathbf{N} . We'll do this by brute force. Here are the steps:

$$[\mathbf{a}^-, \mathbf{N}] = \mathbf{a}^- \mathbf{N} - \mathbf{N} \mathbf{a}^- = \mathbf{a}^- \mathbf{a}^+ \mathbf{a}^- - \mathbf{a}^+ \mathbf{a}^- \mathbf{a}^-.$$

Now, we'll combine the terms in the form

$$[\mathbf{a}^-, \mathbf{N}] = (\mathbf{a}^- \mathbf{a}^+ - \mathbf{a}^+ \mathbf{a}^-) \mathbf{a}^-.$$

This looks complicated until we notice that the expression in the parentheses is just $[\mathbf{a}^-, \mathbf{a}^+]$, which just happens to be 1. Using this fact to simplify, we get

$$[\mathbf{a}^-, \mathbf{N}] = \mathbf{a}^-.$$

We can do the same thing with \mathbf{a}^+ and \mathbf{N} . The result is almost the same except for the sign. Here is the whole list of commutators in one neat package:

$$\begin{aligned}
 [\mathbf{a}^-, \mathbf{a}^+] &= 1 \\
 [\mathbf{a}^-, \mathbf{N}] &= \mathbf{a}^- \\
 [\mathbf{a}^+, \mathbf{N}] &= -\mathbf{a}^+.
 \end{aligned}
 \tag{10.25}$$

This is what you might call a commutator algebra: a set of operators that *closes* under commutation. Commutator algebras have wonderful properties that make them one of the theoretical physicist's favorite tools. We are now going to see the power of this commutator algebra in the iconic example of the harmonic oscillator, using it to find the eigenvalues and eigenvectors of \mathbf{N} . Once we know these, we can immediately read off the eigenvalues of \mathbf{H} from Eq. 10.23. The trick is to use a kind of induction procedure: we begin by supposing we have an eigenvalue and eigenvector of \mathbf{N} . Call the eigenvalue n and the eigenvector $|n\rangle$. By definition,

$$\mathbf{N}|n\rangle = n|n\rangle.$$

Now, let's consider a new vector, obtained by acting on $|n\rangle$ with \mathbf{a}^+ . Let's prove that the result is a different eigenvector of \mathbf{N} , with a different eigenvalue. Again, we accomplish this by straightforward application of the commutation relations. We'll start by writing the expression $\mathbf{N}(\mathbf{a}^+|n\rangle)$ in a slightly more complicated form,

$$\mathbf{N}(\mathbf{a}^+|n\rangle) = [\mathbf{a}^+\mathbf{N} - (\mathbf{a}^+\mathbf{N} - \mathbf{N}\mathbf{a}^+)]|n\rangle.$$

The expression in brackets on the right-hand side is the same as $\mathbf{N}\mathbf{a}^+$, with the term $\mathbf{a}^+\mathbf{N}$ added and then subtracted. But

notice that the expression in parentheses is the last of the commutators from Eqs. 10.25. If we plug that in, we get

$$\mathbf{N}(\mathbf{a}^+|n\rangle) = \mathbf{a}^+(\mathbf{N} + 1)|n\rangle.$$

The last step is to use the fact that $|n\rangle$ is an eigenvector of \mathbf{N} with eigenvalue n . That means we can replace $(\mathbf{N} + 1)$ with $(n + 1)$:

$$\mathbf{N}(\mathbf{a}^+|n\rangle) = (n + 1)(\mathbf{a}^+|n\rangle). \quad (10.26)$$

As always, when we run on autopilot, we have to keep our eyes open for interesting results. Eq. 10.26 is interesting. It says that the vector $\mathbf{a}^+|n\rangle$ is a new eigenvector of \mathbf{N} with eigenvalue $(n + 1)$. In other words, given the eigenvector $|n\rangle$, we have discovered another eigenvector whose eigenvalue is increased by 1. All of this can be summarized by the equation

$$\mathbf{a}^+|n\rangle = |n + 1\rangle. \quad (10.27)$$

Obviously, we can do this again and again to find the eigenvectors $|n + 2\rangle$, $|n + 3\rangle$, and so on. Remarkably, we find that if there is an eigenvalue n , there must be an infinite sequence of eigenvalues above it, spaced by integers. The name *raising* operator seems well chosen.

What about the lowering operator? Not surprisingly, we find that $\mathbf{a}^-|n\rangle$ produces an eigenvector whose eigenvalue is one unit lower:

$$\mathbf{a}^-|n\rangle = |n - 1\rangle. \quad (10.28)$$

This suggests that there must be an unending sequence of eigenvalues below n , but that can't be correct. We already know that the ground state has positive energy, and because $\mathbf{H} = \omega\hbar(\mathbf{N} + 1/2)$ the downward sequence must end. But the only possible way it can end is for there to be an eigenvector $|0\rangle$ such that when \mathbf{a}^- acts on it, the result is zero. (We should not confuse $|0\rangle$ with the zero vector.³) Symbolically, this can be expressed as

$$\mathbf{a}^-|0\rangle = 0. \quad (10.29)$$

Being the lowest energy state, $|0\rangle$ is the ground state, and its energy is $E_0 = \omega\hbar/2$. It is an eigenvector of \mathbf{N} with an eigenvalue 0. We often say that the ground state is *annihilated* by \mathbf{a}^- .

So you see, the abstract construction of \mathbf{a}^+ , \mathbf{a}^- , and \mathbf{N} paid off. It allowed us to find the entire spectrum of harmonic oscillator energy levels without solving a single difficult equation. This spectrum consists of the energy values,

$$\begin{aligned} E_n &= \omega\hbar(n + 1/2) \\ &= \omega\hbar (1/2, 3/2, 5/2, \dots). \end{aligned} \quad (10.30)$$

This quantization of harmonic oscillator energy levels was one of the first results of quantum mechanics, and arguably the most important. The hydrogen atom is a wonderful example of quantum mechanics, but it is, after all, just the

³The 0 vector is the vector whose components are all zero. The vector $|0\rangle$, on the other hand, is a state-vector with nonzero components.

hydrogen atom. The harmonic oscillator, on the other hand, shows up everywhere, from crystal vibrations to electric circuits to electromagnetic waves. The list goes on. Even macroscopic oscillators, like a child on a swing, have quantized energy levels, but the presence of Planck's constant in Eq. 10.30 means that the spacing between levels is so tiny that they are completely undetectable.

The unending spectrum of positive energy levels for a harmonic oscillator is sometimes called a tower, and sometimes called a ladder. It is illustrated schematically in Fig. 10.2.

10.7 Back to Wave Functions

This exercise has amply demonstrated the remarkable power of operator algebras, and the operator method is indeed remarkable. But it's also very abstract. Is it useful in helping us find wave functions, which are more concrete and easier to visualize? Absolutely.

Let's begin with the ground state. We just saw in Eq. 10.29 that the ground state is the unique state that is annihilated by \mathbf{a}^- . Now, let's rewrite Eq. 10.29 in terms of the position and momentum operators, and the ground-state wave function $\psi_0(x)$:

$$\frac{i}{\sqrt{2\omega\hbar}}(\mathbf{P} - i\omega\mathbf{X})\psi_0(x) = 0,$$

or, dividing by the constant factor,

$$(\mathbf{P} - i\omega\mathbf{X})\psi_0(x) = 0.$$

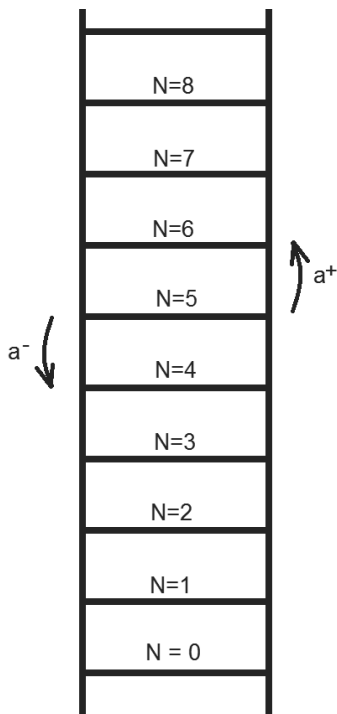


Figure 10.2: Harmonic Oscillator Energy Level Ladder. Energy levels are evenly spaced. \mathbf{a}^+ and \mathbf{a}^- raise and lower the energy level respectively. \mathbf{N} has a lower limit of zero (the ground state), but no upper limit.

If we now replace \mathbf{P} with $-i\hbar \frac{d}{dx}$, we get a first-order differential equation that is much simpler than the second-order Schrödinger equation:

$$\frac{d\psi_0}{dx} = -\frac{\omega x}{\hbar} \psi_0(x).$$

This is a simple differential equation that you can easily

solve. Or, you can just check that the ground-state wave function

$$e^{-\frac{\omega}{2\hbar}x^2}$$

in Eq. 10.15 solves it. Calculating the wave functions for the excited (nonground) states is even easier—we don't even have to solve any equations. Let's go up the ladder to $n = +1$. We can do that by applying \mathbf{a}^+ to the ground state. Let's call the wave function of this new state $\psi_1(x)$.

To avoid dragging the constant $-i/\sqrt{2\omega\hbar}$ around in our calculations, we'll just drop it in our definition of \mathbf{a}^+ . This only affects the numerical coefficient. The resulting equation is

$$\psi_1(x) = (\mathbf{P} + i\omega\mathbf{X})\psi_0(x)$$

or

$$\psi_1(x) = \left(-i\hbar\frac{\partial}{\partial x} + i\omega x\right)e^{-\frac{\omega}{2\hbar}x^2}.$$

Factoring out the i , we get

$$\psi_1(x) = i\left(-\hbar\frac{\partial}{\partial x} + \omega x\right)e^{-\frac{\omega}{2\hbar}x^2}.$$

The “hardest” part of working this out is performing an easy derivative of $e^{-\frac{\omega}{2\hbar}x^2}$. Here is the result:

$$\psi_1(x) = 2i\omega x e^{-\frac{\omega}{2\hbar}x^2},$$

or

$$\psi_1(x) = 2i\omega x\psi_0(x).$$

The only important difference between ψ_0 and ψ_1 is the presence of the factor x in ψ_1 . This has an effect: it causes the wave function of the first excited state to have a zero, or node, at $x = 0$. This is a pattern that continues as we go up the ladder: each successive excited state has an additional node. We can see this pattern emerge by calculating the second excited state at $n = 2$. All we have to do is apply \mathbf{a}^+ again:

$$\psi_2(x) = i\left(-\hbar\frac{\partial}{\partial x} + \omega x\right)\left(xe^{-\frac{\omega}{2\hbar}x^2}\right).$$

We can see right away that the ωx term will result in an ωx^2 term. The $-\frac{\partial}{\partial x}$, meanwhile, will result in two terms because of the product rule for derivatives. One of these terms will come from the exponential (producing another ωx). The other will come from taking the derivative of x . It's clear that what we'll end up with is a quadratic polynomial. If we work out these derivatives, the resulting wave function is

$$\psi_2(x) = (-\hbar + 2\omega x^2)e^{-\frac{\omega}{2\hbar}x^2}.$$

And so it goes, all the way up the ladder. We can see another pattern here: each eigenfunction is a polynomial in x multiplied by $e^{-\frac{\omega}{2\hbar}x^2}$. Because the exponential goes to zero faster

than any of these polynomials grows, each eigenfunction approaches zero asymptotically as x goes to plus or minus infinity. Also, because the degree of each polynomial is one greater than the degree of the previous one, each eigenfunction has one more zero than the previous one.⁴ This also explains why successive eigenfunctions alternate between being symmetric and antisymmetric. Specifically, eigenfunctions with polynomials of even degree are symmetric, while those with polynomials of odd degree are antisymmetric. The polynomials in this sequence are very well-known. They're called the *Hermite* polynomials. The ground-state eigenfunction $e^{-\frac{\omega}{2}x^2}$, which appears in all of these higher-energy eigenfunctions, is symmetric in x .

Fig. 10.3 displays the eigenfunctions for several different energy levels. Each successive eigenfunction oscillates more rapidly than the one before it. This corresponds to an increase in momentum. The more rapidly the wave function oscillates, the greater the momentum of the system. At higher energy levels, the wave function also becomes more spread out. In physical terms, this means the mass is moving farther from the equilibrium point, and moving faster.

These eigenfunctions contain another important lesson. Although they approach zero asymptotically (quite rapidly) they never quite reach zero. That means there is a small but finite chance of finding the particle “outside the bowl” that defines its potential energy function. This phenomenon,

⁴It turns out that these zeros occur for real values of x , but that's not obvious from what we've seen. In a physical sense, the zeros seem a little weird, because they are points where the moving mass will never be found, even though it's merrily whizzing back and forth.

known as *quantum tunneling*, is completely unknown in classical physics.

10.8 The Importance of Quantization

We've climbed a high mountain in these lectures, but it's not the last mountain. Looking out from the present vantage point, we can get a glimpse of the enormous landscape of quantum field theory. That's material for another book. Or maybe three. But still, we can see a bit of the terrain from where we are.

Consider the example of electromagnetic radiation in a cavity, as shown in Fig. 10.4. In this context, a cavity is a region of space bracketed by a pair of perfectly reflecting mirrors that keep the radiation bouncing endlessly back and forth. Think of the cavity as a long metallic tube that the radiation can travel along in both directions.

There are many wavelengths that can fit into the cavity. Let's consider waves of length λ . Like all waves, these waves oscillate, very much like a mass on the end of a spring. But it's important not to get confused here: the oscillators are not masses attached to springs. What's really oscillating are the electric and magnetic fields. For each wavelength, there is a mathematical harmonic oscillator describing the amplitude or strength of the field. That's a lot of harmonic oscillators all running simultaneously. Fortunately, however, they all oscillate independently, so we can focus our attention on waves of one particular wavelength and ignore all the

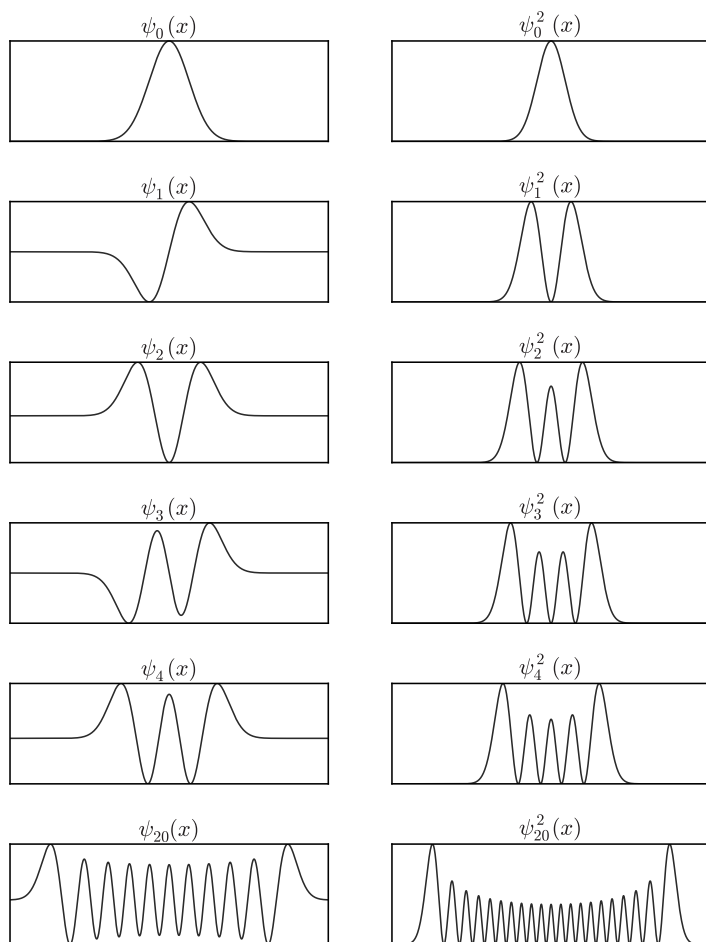


Figure 10.3: Harmonic Oscillator Eigenfunctions. Amplitudes are shown on the left, probabilities on the right. The higher-energy wave functions oscillate more rapidly and are more spread out.

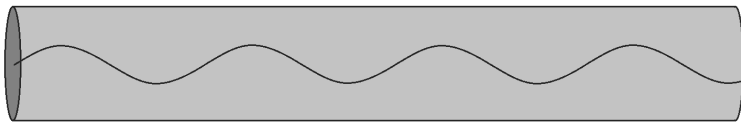


Figure 10.4: Electromagnetic Radiation in a Cavity

others.

There is only one important number associated with a harmonic oscillator—namely, its frequency. You probably already know how to calculate the frequency of a wave of length λ :

$$\omega = \frac{2\pi c}{\lambda}.$$

In classical physics, of course, the frequency is just the frequency. But in quantum mechanics, the frequency determines the quantum of energy of the oscillator. In other words, the energy contained in waves of length λ has to be

$$(n + 1/2)\hbar\omega.$$

The term $(1/2)\hbar\omega$ is not important for our purposes. It's called the zero-point energy, and we can ignore it. If we do, the energy of waves of length λ becomes

$$\frac{2\pi\hbar c}{\lambda}n,$$

where n can be any integer from zero on up. In other words,

the energy of an electromagnetic wave is quantized in indivisible units of

$$\frac{2\pi\hbar c}{\lambda}.$$

For a classical physicist this is very odd. No matter what you do, the energy always comes in unbreakable units.

You may already know that these units are called photons. In fact, *photon* is just another name for the quantized unit of energy in a quantum harmonic oscillator. But we can also describe the same facts another way. Being indivisible, photons can be thought of as elementary particles. A wave excited to its n th quantum state can be thought of as a collection of n photons.


What is the energy of a single photon? That's easy. It's just the energy that it takes to add one more unit, namely

$$E(\lambda) = \frac{2\pi\hbar c}{\lambda}.$$

Here, we can see something that has dominated physics for well over a century: the shorter the wavelength of a photon, the higher its energy. Why would a physicist be interested in making short-wavelength photons, given that they are costly in energy? The answer is to see more clearly. As discussed in Lecture 1, to resolve an object of a given size, you must use waves of that size or smaller. To see a human figure, a wavelength of a few inches is good enough. To see a tiny speck of dust, you may need visible light of a much smaller wavelength. To resolve the parts of a proton, the wavelength must be smaller than 10^{-15} meters, and the corresponding

photons must be very energetic. In the end, it all goes back to the harmonic oscillator.

On that note, my friends, we conclude this volume of the Theoretical Minimum series. I look forward to seeing you in Special Relativity.



© Margaret Sloan

Appendix

Pauli Matrices

$$\sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

$$\sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}$$

Action of Spin Operators

$$\begin{aligned}
 |u\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix} & \iff \begin{aligned} \sigma_z|u\rangle &= |u\rangle \\ \sigma_x|u\rangle &= |d\rangle \\ \sigma_y|u\rangle &= i|d\rangle \end{aligned} \\
 |d\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix} & \iff \begin{aligned} \sigma_z|d\rangle &= -|d\rangle \\ \sigma_x|d\rangle &= |u\rangle \\ \sigma_y|d\rangle &= -i|u\rangle \end{aligned} \\
 |r\rangle = \begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix} & \iff \begin{aligned} \sigma_z|r\rangle &= |l\rangle \\ \sigma_x|r\rangle &= |r\rangle \\ \sigma_y|r\rangle &= -i|l\rangle \end{aligned} \\
 |l\rangle = \begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{-1}{\sqrt{2}} \end{pmatrix} & \iff \begin{aligned} \sigma_z|l\rangle &= |r\rangle \\ \sigma_x|l\rangle &= -|l\rangle \\ \sigma_y|l\rangle &= i|r\rangle \end{aligned} \\
 |i\rangle = \begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{i}{\sqrt{2}} \end{pmatrix} & \iff \begin{aligned} \sigma_z|i\rangle &= |o\rangle \\ \sigma_x|i\rangle &= i|o\rangle \\ \sigma_y|i\rangle &= |i\rangle \end{aligned} \\
 |o\rangle = \begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{-i}{\sqrt{2}} \end{pmatrix} & \iff \begin{aligned} \sigma_z|o\rangle &= |i\rangle \\ \sigma_x|o\rangle &= -i|i\rangle \\ \sigma_y|o\rangle &= -|o\rangle \end{aligned}
 \end{aligned}$$

Change of Basis

$$|r\rangle = \frac{1}{\sqrt{2}}|u\rangle + \frac{1}{\sqrt{2}}|d\rangle$$

$$|l\rangle = \frac{1}{\sqrt{2}}|u\rangle - \frac{1}{\sqrt{2}}|d\rangle$$

$$|i\rangle = \frac{1}{\sqrt{2}}|u\rangle + \frac{i}{\sqrt{2}}|d\rangle$$

$$|o\rangle = \frac{1}{\sqrt{2}}|u\rangle - \frac{i}{\sqrt{2}}|d\rangle$$

Spin Component in the \hat{n} Direction

Vector Notation

$$\sigma_n = \vec{\sigma} \cdot \hat{n}$$

Component Form

$$\sigma_n = \sigma_x n_x + \sigma_y n_y + \sigma_z n_z$$

More Concretely

$$\sigma_n = n_x \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} + n_y \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} + n_z \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

Combined in a Single Matrix

$$\sigma_n = \begin{pmatrix} n_z & (n_x - in_y) \\ (n_x + in_y) & -n_z \end{pmatrix}$$

Spin Operator Multiplication Tables

A word about notation: Table 3 below uses the symbol i in two different ways. Inside a ket, such as $|io\rangle$, it is part of a state-label— io signifies “*in-out*.” But when i appears outside of a ket symbol, as in $i|oo\rangle$, it signifies the unit imaginary number.

Table 1: *Up-Down* Basis

	2-Spin Eigenvectors			
	$ \mathbf{uu}\rangle$	$ \mathbf{ud}\rangle$	$ \mathbf{du}\rangle$	$ \mathbf{dd}\rangle$
σ_z	$ uu\rangle$	$ ud\rangle$	$- du\rangle$	$- dd\rangle$
σ_x	$ du\rangle$	$ dd\rangle$	$ uu\rangle$	$ ud\rangle$
σ_y	$i du\rangle$	$i dd\rangle$	$-i uu\rangle$	$-i ud\rangle$
τ_z	$ uu\rangle$	$- ud\rangle$	$ du\rangle$	$- dd\rangle$
τ_x	$ ud\rangle$	$ uu\rangle$	$ dd\rangle$	$ du\rangle$
τ_y	$i ud\rangle$	$-i uu\rangle$	$i dd\rangle$	$-i du\rangle$

Table 2: *Right-Left* Basis

2-Spin Eigenvectors				
	$ \mathbf{rr}\rangle$	$ \mathbf{rl}\rangle$	$ \mathbf{lr}\rangle$	$ \mathbf{ll}\rangle$
σ_z	$ lr\rangle$	$ ll\rangle$	$ rr\rangle$	$ rl\rangle$
σ_x	$ rr\rangle$	$ rl\rangle$	$- lr\rangle$	$- ll\rangle$
σ_y	$-i lr\rangle$	$-i ll\rangle$	$i rr\rangle$	$i rl\rangle$
τ_z	$ rl\rangle$	$ rr\rangle$	$ ll\rangle$	$ lr\rangle$
τ_x	$ rr\rangle$	$- rl\rangle$	$ lr\rangle$	$- ll\rangle$
τ_y	$-i rl\rangle$	$i rr\rangle$	$-i ll\rangle$	$i lr\rangle$

Table 3: *In-Out* Basis

2-Spin Eigenvectors				
	$ \mathbf{ii}\rangle$	$ \mathbf{io}\rangle$	$ \mathbf{oi}\rangle$	$ \mathbf{oo}\rangle$
σ_z	$ oi\rangle$	$ oo\rangle$	$ ii\rangle$	$ io\rangle$
σ_x	$i oi\rangle$	$i oo\rangle$	$- ii\rangle$	$- io\rangle$
σ_y	$ ii\rangle$	$ io\rangle$	$- oi\rangle$	$- oo\rangle$
τ_z	$ io\rangle$	$ ii\rangle$	$ oo\rangle$	$ oi\rangle$
τ_x	$i io\rangle$	$-i ii\rangle$	$i oo\rangle$	$-i oi\rangle$
τ_y	$ ii\rangle$	$- io\rangle$	$ oi\rangle$	$- oo\rangle$

Index

- 2 × 2 matrices, combining, 188
- 3-vector operators, 75, 83–85, 119
- 3-vectors, 25, 27, 74–75, 83
 - orthogonal unit vectors and, 32–33
- 4 × 4 matrices, from combined 2 × 2 matrices, 188
- Addition
 - of complex numbers, 23
 - vector, 26
- Amplitude, 39, 108, 342, 343
 - for paths, 306–309
- and* rule, 14, 15, 20
- Annihilation operators, 327–337
- Anti-Hermitian operator, 250
- Antisymmetric eigenfunctions, 341
- Apparatus, measurement and, 5–13,
37–38, 71, 75, 81–82, 83–84, 91,
126–127, 180, 219–224, 227–230
- Associative property, 26, 193, 239
- Atoms, 259, 290, 311
 - in crystal lattice, 313
 - hydrogen, 336–337
 - quantum mechanics and, 2, 71,
149, 316
 - size of, 104
 - spins of, 180–181
 - wave packets and, 297, 301
- Average, 140–141, 157–158, 213, 271,
286, 288, 292, 295
 - bra-ket notation for, 106–107
 - defining, 105–106
 - See also* Expectation values
- Average value, 105
- Axioms, vector space, 24–27
- Basis of simultaneous eigenvectors,
131–133
- Basic vectors, 32–34, 38, 40, 41,
48–49, 54, 55, 64, 67, 97,
98, 106, 120–125, 130–136,
173, 185, 189, 191, 195, 196,
198, 202, 204, 208, 210, 211,
219, 224, 236, 237, 251, 258,
260–263, 275
 - components, 56
 - entangled states, 165–167
 - labeling, 150.151, 152, 153, 154,
160–163
 - product states, 163–165
- Bell, John, 223, 227
- Bell's theorem, 227–231
- Boolean logic, 13–18
- Bracket notation, 11
- Bra-ket notation, 105
 - for averages, 106–107
- Bras (bra vectors), 28–30, 240
 - inner product and, 30–32
 - linear operators and, 58–59
 - outer products and, 194

- Canonical momentum, 315, 318–320
- Canonical momentum conjugate to x , 315, 318–320
- Cartesian coordinates, 89, 116, 136
- Cartesian representation, of complex number, 22
- Cauchy-Schwarz inequality, 142
 - triangle inequality and, 142–146
- Change
 - in classical physics, 94
 - continuity and, 100
 - unitarity and incremental, 100
- Classical entanglement, 155–160
- Classical equations, quantization and, 289–290
- Classical limit, 295–301
- Classical physics
 - change in, 94
 - change in expectation values over time and, 109–114
 - commutators and, 266–268
 - momentum in, 255
 - particle dynamics and, 279
 - pure and mixed states and, 199–200
 - quantum mechanics vs., 2–3
 - testing propositions of, 16–18
- Collapse of the wave function, 126–127
- Column vectors, 27–28, 49
 - kets and, 29
 - spin states as, 47
- Commutation relations, 118, 119, 138–139, 287, 309, 328, 332, 334
- Commutative property, 26
- Commutator algebra, 334–337
- Commutators, 111–116, 138, 142, 146, 147, 269, 280, 287, 293, 294,
 - classical physics and, 266–268
 - operators and, 328, 330, 332, 333, 334, 335
 - Poisson brackets and, 112–114, 265–268
- Commuting variables, complete sets of, 129–136
 - wave functions, 134–136
- Complex conjugate, 23
- Complex conjugate numbers, 28, 30
- Complex conjugation, for operators, 59–61
- Complex numbers, 21–30, 34, 38, 42, 44
 - addition of, 23
 - eigenvalues and, 58
 - multiplication of, 23
 - phase-factors, 24
 - representations of, 22
- Complex vector spaces, orthonormal basis and, 33
- Component matrices, building tensor product matrices from, 188–192
- Component, 56
 - of 3-vector, 25, 74–75, 83, 116
 - addition of, 27
 - of angular momentum, 119
 - of basis vector, 56
 - of generic state, 38
 - inner products and, 31, 34
 - multiplication of, 28
 - of phase factor, 24
 - of spin, 9, 13, 16–17, 20, 37, 69, 71, 75, 77, 83–84, 87, 90–91, 116–117, 119, 130–131, 138–139, 162, 167–168, 170, 174–175, 176, 178–179, 180–181, 218, 222, 251, 257, 260, 349
 - of spin operator, 71–72, 75, 116
 - of state-vector, 40, 227, 237, 336n
 - of system, 154, 222
 - of vector, 8, 9–10
 - wave functions and, 136
- Component form
 - of addition, 23, 27–28
 - of bra-vectors, 59
 - equation in, 54, 59, 79
 - of multiplication, 58–59
 - of tensor product operators, 155, 171–172, 184, 188, 204

- Component matrices, 188–192
- Composite observables, 175–181
- Composite operator
 - composite vectors and, 171
 - energy and measurement of, 180–181
- Composite state, two spin, 161–181
- Composite systems
 - mixed and pure states and, 200–201
 - observables in, 167–175
 - product states, 163–165
 - representing, 151–155
 - tensor products and, 150–155
 - See also* Entanglement
- Composite vectors, composite operators and, 171
- Conservation of distinctions, 97–99
- Conservation of energy, 114–115
- Conservation of overlaps, 99
- Continuity, 100–101
- Continuous functions, 236–250
 - functions as vectors, 238–245
 - integration by parts, 245–246
 - linear operators, 246–250
 - wave functions and, 236–238
- Correlation
 - of near-singlet state, 234
 - of product state, 232
 - of singlet state, 233
- Correlation test for entanglement, 213–214
- Creation operators, 327–337
- Crystal lattice, atom in, 313
- Degeneracy, 64
- Density matrices, 184, 196–199
 - calculating, 210–212
 - entanglement and, 199–202
 - of near-singlet state, 234
 - notation for, 201–202
 - of product state, 232
 - properties of, 207
 - for single spin, 202–203
 - of singlet state, 233
 - two-spin system and, 203–217, 231
- Density matrix test for entanglement, 214–218
- Determinism
 - in classical physics, 94
 - in quantum mechanics, 9–11, 96
- Dirac, Paul, 105, 113, 194, 278
- Dirac delta functions, 241, 242–245, 253
- Dirac's bracket notation, 11
- Distributive property, 26
- Dot product, 30, 31, 144, 180
- Down states, 219–221
- Dual number systems, 23
- Eigen-equation, 256
- Eigenfunctions, 253
 - alternation between being symmetric and antisymmetric, 340–341
 - for energy levels, 341, 343
 - harmonic oscillator, 341, 343
- Eigenstate, collapse of the wave function and, 126–127
- Eigenvalues, 56–59, 70, 71–72
 - of density matrix, 207, 215–217
 - energy, 121, 322–323
 - of Hermitian operators, 62–63
 - of operators, 80
 - of position, 252–254
 - of spin operator, 76, 77–78
- Eigenvectors, 56–59, 70
 - of annihilation operator, 328
 - of creation operator, 328
 - defined, 57
 - energy, 121, 322–323
 - of Hermitian operator, 64–67
 - of momentum, 255–260
 - of operators, 80
 - of position, 252–254
 - of projection operator, 194
 - simultaneous, 131–133
 - of spin operator, 76, 77–80
- Einstein, Albert, 155, 175, 223, 227
- Electric current, 313

- Electromagnetic radiation in cavity, 342–345
- Electromagnetic waves, 313
- Electrons, 2, 149, 259, 301
 spin of, 3–4, 116, 180, 290
 wave packets and, 301
 waves and, 235
- Energy
 composite operator and, 180–181
 conservation of, 114–115
 creation and annihilation operators and, 328–337
 frequency and, 123
 harmonic oscillator and, 314–316, 317–319
 of particle with negative momentum, 278
 of photon, 345
See also Hamiltonian
- Energy eigenvalues, 121, 322–323
- Energy eigenvectors, 121, 322–323
- Energy levels
 eigenfunctions for, 341, 343
 harmonic oscillators and, 322–323, 336–337, 338
- Entangled states, 165–167
- Entanglement, 149–181
 Bell's Theorem and, 227–231
 classical, 155–160
 combining quantum systems, 160–161
 composite observables, 175–181
 correlation test for, 213–214
 density matrices and, 184, 199–202, 210–212
 density matrix test for, 214–218
 entangled states, 165–167
 example: calculating a density matrix, 210–212
 locality and, 223–226
 of near-singlet state, 234
 observables and, 167–175
 process of measurement and, 218–223
 of product state, 163–165, 232
 of singlet state, 233
 summary of, 231–234
 tests for, 212–218
 for two spins, 161–163, 202–210
- Euler-Lagrange equations, 305n
- Expectation values, 87–88, 91, 105–108
 change over time in, 109–114
 conservation of, 115
 correlation test for entanglement and, 213–214
 for density matrix, 198
 of entangled state, 172–175
 of near-singlet state, 234
 particle dynamics and, 278–279
 of product state, 232
 of projection operator, 195–196
 of singlet state, 233
 in spin over time, 116–119
- Experiments
 apparatus and, 5–13
 invasiveness of, 12–13
 probabilities for outcomes of (*see* Probabilities for experimental outcomes)
 two-state system, 4–11
- Feynman, Richard, 302, 309
- Forces, 290–294
- Fourier transforms, 260–261, 265, 285
- Frequency
 energy and, 123
 of harmonic oscillator, 344–345
- Functions
 Dirac delta, 241, 242–245, 253
 Gaussian, 327
 normalizable, 318
 potential, 291, 297–298
 probability, 105–106, 213, 295
 as vectors, 238–245
 vector space, 27–28
 zero, 239
See also Continuous functions; Eigenfunctions; Wave functions
- Fundamental theorem of quantum mechanics, 64

- Gaussian curve, 301
- Gaussian function, 327
- Gaussian wave packets, 301
- General Schrödinger equation, 102, 274
- General uncertainty principle, 146–148, 268, 269–270
- Gluons, 259
- Gram-Schmidt procedure, 67–69
- Gravitons, 280
- Ground states, 324–327
 - annihilation of, 336
 - wave functions for, 337–339
- Hamiltonian, 99–102
 - canonical momentum and, 319–320
 - conservation of, 115
 - entanglement and, 181
 - for harmonic oscillator, 318–320, 321, 322–323, 324–326, 329–334, 336
 - motion of particles and, 274–278
 - nonrelativistic free particles and, 280–283
 - quantum, 101, 103
 - spin in magnetic field, 116–119
 - time evolution of system and, 274
- Hamiltonian operator, Schrödinger ket and, 124
- Hamilton's equations, 274, 279
- Harmonic oscillator, 311–346
 - annihilation operators, 327–337
 - classical description, 314–316
 - creation operators, 327–337
 - energy levels, 322–323
 - ground state, 324–327
 - prevalence in physics, 311–313
 - quantization and, 342–346
 - quantum mechanical description, 316–321
 - Schrödinger equation, 321–322
 - wave functions, 337–342
- Harmonic oscillator energy level
 - ladder/tower, 337, 338
- Heisenberg, Werner, 327
- Heisenberg Uncertainty Principle, 139–140, 148, 269–271
- Hermite, Charles, 62
- Hermite polynomials, 341
- Hermitian
 - density matrices as, 207, 208
 - momentum as, 262
 - position as, 262
 - projection operators as, 194
- Hermitian conjugation/conjugate, 59–61, 62, 63, 65, 97–98, 100, 332
- Hermitian matrix, 62, 137–138, 195n, 208
- Hermitian observable, 262
- Hermitian operators, 52, 101, 112, 138, 255
 - action on state-vector, 107–108
 - in composite space of states, 168
 - eigenvector of, 139, 236, 262
 - expectation value of, 109
 - linear operators as, 70, 73–74, 246–250
 - orthonormal bases and, 64–67
 - orthonormal edge vectors of, 136
 - overview, 61–63
 - particles and, 252
 - trace of, 196
- Hilbert, David, 239
- Hilbert spaces, 25, 239
- Hooke's law, 312
- Hydrogen atom, 336–337
- Identity, resolving, 261–264
- Identity operator, from projection operators, 195
- Inner products, 28–29, 30–32, 193
- Integrals, replacing sums, 240, 241
- Integration by parts, 245–246
- Kets (ket vectors), 28–30
 - axioms of, 25–27
 - composite systems and, 153–154
 - inner product, 30–32
 - Schrödinger, 124–126
- Kinematics, 317

- Kronecker delta, 205
 replaced by Dirac delta functions,
 241, 242–245
- Kronecker product, 188–192, 205n
- Kronecker symbol, 98, 161
- Lagrange equation, 314–316
- Lagrangian, 302–303, 314–316, 318,
 319
- Law of evolution, 5
- Least action principle, 301–305
- Linearity, 27, 53
- Linear motion, 295–301
- Linear operators, 52–69, 246–250
 eigenvalues, 56–59
 eigenvectors, 56–59
 Gram-Schmidt procedure, 67–69
 Hermitian conjugation, 59–61
 Hermitian operators, 61–63
 Hermitian operators, orthonormal
 bases and, 64–67
 machines and matrices, 52–56
 observables and, 69–70, 73
 outer product as, 193–196
 properties of, 53
 time-development operator, 97
- Liouville's theorem, 274
- Locality
 defined, 223–224
 Einstein vs. Bell and, 227
 entanglement and, 223–226
- Lowering operators (annihilation
 operators), 327–337
- Machines, matrices and, 52–56
- Magnetic field, spin in, 116–119
- Mathematical concepts
 complete sets of commuting
 variables, 129–136
 complex numbers, 21–24
 continuous functions, 236–250
 functions as vectors, 238–245
 integration by parts, 245–246
 linear operators, 52–69, 246–250
 outer products, 193–196
 tensor products, 149–155
 tensor products in component
 form, 184–192
 vector spaces, 24–34
- Matrices
 $4 \cdot 4$, 188
 machines and, 52–56
 Pauli, 80, 118, 137
 tensor product, building, 185–192
 $2 \cdot 2$, 188 [is this entry out of
 lexical sequence?]
- Matrix elements, 55
- Matrix multiplication, 56, 59
- Matrix notation, transposing in,
 60–61
- Maximally entangled state, 217, 221
- Maxwell's equations, 290
- Mean value, 105
- Measurables, states that depend on
 more than one, 129–133
- Measurement, 137–139
 apparatus and, 5–11, 219–223
 collapse of the wave function and,
 126–127
 multiple, 129–133
 operators and, 80–82
 process of, 218–223
 states and, 2–3
- Minimum-uncertainty wave packets,
 301
- Minus first law, 94, 274
 quantum version of, 94–95, 97
- Mixed states, 198, 199–200
 composite system and, 200–201
 density matrices and, 208–209
- Momentum
 canonical, 315, 318–320
 connection between quantum and
 classical physics, 268
 eigenfunctions and, 341
 eigenvectors of, 255–260
 forces and, 292–294
 Heisenberg Uncertainty Principle
 and, 269
 proposition for, 20–21
 velocity and, 286–288, 293
 wavelength and, 259–260

- Momentum basis, 260–265
- Momentum operator, 255–257
- Momentum representation, of wave function, 260–265
- Motion of particles. *See* Particle dynamics
- Multiplication
 - of column vector, 28
 - of complex numbers, 23
 - matrix, 56, 59
 - vector, 26
- Near-singlet state
 - correlation, 234
 - density matrix, 234
 - description of, 234
 - entanglement status of, 234
 - expectation values, 234
 - normalization, 234
 - state-vector, 234
 - wave function, 234
- Negation, 14
- Neutrino, 3
 - moving at speed of light, 277–278
- Newton's law, 291, 292
 - quantum version of, 293–294
- Nonlocality, 231
- Nonrelativistic free particles, 280–283
- Normalizable functions, 318
- Normalization
 - of near-singlet state, 234
 - of product state, 232
 - of singlet state, 233
- Normalized vector, 32, 40
- not* rule, 14
- Number operator, 332–333
- Observables
 - complete set of commuting, 133
 - composite, 175–181
 - composite system, 167–175
 - defined, 52
 - linear operators and, 69–70, 73
 - multiple, 130–131
- Observations, collapse of the wave function and, 126–127
- Operator method
 - harmonic oscillator and, 328–337
 - wave functions and, 337–342
- Operators
 - 3-vector, 75, 83–85, 119
 - annihilation, 327–337
 - anti-Hermitian, 250
 - commutators and, 328, 334
 - composite, 171, 180–181
 - creation, 327–337
 - Hamiltonian, 124
 - Hermitian (*see* Hermitian operators)
 - identity, 195
 - linear (*see* Linear operators)
 - measurement and, 80–82
 - misconception regarding, 81–82
 - momentum, 255–257
 - number, 332–333
 - projection, 194–195
 - spin, 74–80
 - state-vectors and, 80–81
 - time-development, 95, 97–99
 - time-evolution, 99–102
 - unitary, 95, 97–99
 - zero, 133
- Original Schrödinger equation, 274
 - nonrelativistic free particle and, 281–283
- or* rule, 14, 15, 19
- Orthogonal basis vectors, 48
- Orthogonal states, 39–40, 97
- Orthogonal state-vectors, 70, 72
- Orthogonal vectors, 32, 64–67, 70
- Orthonormal bases, 32–34
 - Gram-Schmidt procedure, 67–69
 - Hermitian operators and, 64–67
- Outer products, 193–196
- Overlap, 72, 73
- Parameters, counting, 45–47
- Partial derivatives, time and, 320–321

- Particle dynamics, 273–309
 - example, 273–279
 - forces, 290–294
 - linear motion and classical limit, 295–301
 - nonrelativistic free particles, 280–283
 - path integrals, 301–309
 - quantization, 288–290
 - time-independent Schrödinger equation, 283–285
 - velocity and momentum, 286–288
- Particle moving in three-dimensional space, measuring, 130
- Particles, 235–236
 - coordinates of, 238
 - Heisenberg Uncertainty Principle and, 269–271
 - Hermitian operators and, 252
 - wave function and probability for finding position of, 260–265
- Particles, state of, 250–260
 - eigenvalues and eigenvectors of position, 252–254
 - momentum and its eigenvectors, 255–260
- Particle-wave duality, 236
- Path integrals, 301–309
- Pauli matrices, 80, 118, 137
- Phase ambiguity, 42
- Phase-factors, 24, 42, 46, 108–109
- Phase indifference, 47, 48–49
- Photons, 260, 277, 280, 345–346
- Planck's constant, 102–104, 148, 255, 337
- Poisson brackets, 280
 - commutators and, 112–114, 265–268
- Polarization vector, 91
- Polar representation of complex number, 22–23
- Position
 - eigenvalues and eigenvectors of, 252–254
 - Heisenberg Uncertainty Principle and, 269
 - proposition for, 20–21
 - Position representation, of wave function, 260–262, 263–265
- Potential functions, 291
 - spiky, 297–298
- Precession, of spin in magnetic field, 119
- Principle of Least Action, 301–305
- Principle of Stationary Action, 302n
- Probabilities for experimental outcomes, 8, 19, 48–49, 70, 72–73, 87–90, 238, 306
 - replaced by probability densities, 241, 242
 - Schrödinger ket and, 124–126
- Probability
 - entanglement and, 206–207, 222
 - wave function and, 260–261, 264, 270
- Probability amplitudes, 39, 108–109
- Probability density, 199, 317, 325
 - replacing probabilities, 241, 242
- Probability distribution, 110, 112, 213
 - in classical mechanics, 158–159
 - particle dynamics and, 278–279
 - uncertainty and, 140–141
- Probability function, 105–106, 213, 295
- Product states, 163–165
 - correlation, 232
 - counting parameters for, 165
 - density matrix, 232
 - density matrix test for entanglement and, 215–218
 - description of, 232
 - entanglement status, 232
 - expectation values, 232
 - normalization, 232
 - state-vector, 232
 - wave function, 232
- Projection operators, 194
 - properties of, 194–195
- Propositions
 - classical, 13–16
 - classical, testing, 16–18
 - quantum, testing, 18–21

- Pure states, 198, 199–200
 - composite system and, 200–201
 - density matrices and, 207–209, 217
- Quantization, 288–290
 - importance of, 342–346
- Quantum abstractions, 2
- Quantum electrodynamics, 290
- Quantum field theory, 342
 - path integrals and, 309
- Quantum Hamiltonian, 101, 103
- Quantum mechanics
 - as calculus of probabilities, 36
 - classical mechanics *vs.*, 2–3
 - conservation of energy and, 114–115
 - focus of, 1–3
 - fundamental theorem of, 64
 - Planck’s constant and, 102–104
 - testing propositions of, 18–21
- Quantum mechanics, principles of, 69–74, 99
 - 3-vector operators, 83–85
 - application, 85–90
 - operators, measurement and, 80–82
 - spin operators, 74–75
 - spin operators, constructing, 75–80
 - spin-polarization principle, 90–91
- Quantum Sim, 227–231
- Quantum spins, 3–4, 36–37, 227, 229–230
- Quantum states, 35–49
 - along the x axis, 41–42
 - along the y axis, 42–45
 - counting parameters, 45–47
 - incompleteness of, 36
 - representing spin states as column vectors, 47
 - spin states, 37–40
 - states and vectors, 35–37
- Quantum systems, combining, 160–161
- Quantum tunneling, 341–342
- Quarks, 3, 259, 311
- Qubits, 3–4, 5, 311
 - measuring system of two, 130–131
- Raising operators (creation operators), 327–337
- Real numbers, quantum mechanics and, 61–63
- Reversibility, 94
- Row vectors, bras and, 29–30
- Schrödinger, Erwin, 327
- Schrödinger equations
 - generalized (*see* Time-dependent Schrödinger equation)
 - original, 274, 281–283
 - path integrals and, 309
 - solving, 119–124
 - spin state evolution and, 227–230
 - time-dependent (*see* Time-dependent Schrödinger equation)
 - for time derivatives, 110–112
 - time-independent, 120–121, 124, 283–285, 286, 289
- Schrödinger ket, 124–126
- Schrödinger’s Ket, 102
- Sets, Boolean logic and, 13–16
- Simultaneous eigenvectors, 131–133
- Singlet state, 166–167, 181
 - correlation, 233
 - density matrix, 233
 - description of, 233
 - entanglement status of, 233
 - expectation values, 233
 - normalization, 233
 - state-vector, 233
 - wave function, 233
- Space of states, 4–5, 13, 16, 24, 25, 37, 40, 44, 71, 94, 124, 150–151, 160, 162, 165, 166, 167–168, 216, 219, 238, 274, 289, 317
- Speed of light, particles moving at, 277–278
- Spherical coordinates, 89–90

- Spin
- 3-vector operators and, 83–85
 - along the x axis, 41–42
 - along the y axis, 42–45
 - density matrix for, 202–203
 - expectation values of, 87–88, 91
 - interaction with apparatus, 5–13
 - in magnetic field, 116–119
 - number of distinct states for, 45–47
 - quantum, 3–4, 36–37, 227, 229–340
 - uncertainty principle and, 20
- See also* Qubits; Two spins
- Spin components, simultaneous measurement of, 138–139
- Spin operators, 74–75
- constructing, 75–80
- Spin-Polarization Principle, 90–91, 172
- Spin states
- as column vectors, 47
 - representing, 37–40
 - Schrödinger's equation and evolution of, 227–230
- Spring constant, 312
- Standard deviation, 140, 141
- State
- of apparatus, 219–220
 - change over time, 94, 274
 - maximally entangled, 217, 221
 - measurement and, 2–3
 - mixed (*see* Mixed states)
 - near-singlet (*see* Near-singlet state)
 - of particles, 250–260
 - pure (*see* Pure states)
 - quantum (*see* Quantum states)
 - in quantum mechanics, 2
 - singlet (*see* Singlet state)
 - that depend on more than one measurable, 129–133
 - triplet, 166–167, 179, 181
 - unambiguously distinct, 70, 72
- State-labels, for composite system, 152, 153, 154, 160–161
- State of system, in classical vs. quantum physics, 21, 273–274
- State space, Boolean logic and, 13–16
- State-vectors, 70
- action of Hermitian operator on, 107–108
 - as complete description of system, 175
 - evolution of with time, 99
 - of near-singlet state, 234
 - operators and, 80–81
 - phase-factor and, 108–109
 - physical properties of, 46
 - of product state, 163–165, 232
 - representing spin states using, 37–40
 - of singlet state, 233
 - time derivative of, 102
 - time evolution of, 95–96
 - wave functions and, 136
- See also* Bras (bra vectors); Kets (ket vectors); Singlet state; Triplet states
- Statistical correlation, 158
- Subset, 13, 14, 15–16
- Sums, integrals replacing, 240
- Symmetric eigenfunctions, 340–341
- Systems
- number of parameters characterizing, 45–47
 - quantum, combining, 160–161
- See also* Composite systems; Two-spin system
- Tensor products, 149–155, 165, 176
- Tensor products in composite form, 184–192
- building tensor product matrices from basic principles, 185–187
 - building tensor product matrices from component matrices, 188–192

- Tests for entanglement, 212–218
- Time
 - change in expectation values over, 109–114
 - conservation of distinctions and, 97–99
 - determinism and, 96
 - partial derivatives and, 320–321
 - time-evolution operator, 99–102
 - unitarity, 95, 98–99
 - See also* Schrödinger equations
- Time dependence, 116, 125, 286, 322. *See also* Uncertainty
- Time-dependent Schrödinger equation, 102
 - harmonic oscillation and, 321–323
 - particle dynamics and, 274, 275–276, 289, 302
 - solving, 120, 121–124
 - state of system and, 126
- Time derivatives, 102
 - Schrödinger equation for, 110–112
- Time-development operator, 95
 - conservation of distinctions and, 97–99
- Time evolution, 274
 - determinism and, 96
 - entanglement and, 181
 - unitary operators and, 98–99
- Time-evolution operator, 99–102
- Time-independent Schrödinger equation, 120–121, 124
 - particle dynamics and, 283–285, 286, 289
- Trace
 - of density matrix, 206, 207, 209
 - of projection operator, 195, 196
 - properties of, 209
- Trajectories, path integrals, 301–309
- Transposing, 60–61
- Triangle inequality, 142–146
- Triplet states, 166–167, 179, 181
- Truth-value, 13–14
- Two spins, 161–181
 - entanglement for, 202–210
- Two-spin system
 - Bell's theorem and, 230–231
 - density matrix of, 202–212, 214–218, 226, 231
- Two-state system, experiment on, 4–11
- Uncertainty
 - Cauchy-Schwarz inequality, 142
 - defined, 140–141
 - triangle inequality and Cauchy-Schwarz inequality, 142–146
- Uncertainty principle, 20, 139–140, 146–148
 - Heisenberg, 139–140, 148
- Unitarity, 95, 98–99, 100
- Unitary evolution, 218, 222, 225
- Unitary matrix, 225
- Unitary operators, 95, 97–99
- Unitary time evolution, 181
- Unit matrix, 137
 - density matrix and, 217
- Unit (normalized) vector, 32
 - state of system and, 40
- Unit operator, as observable, 138
- Up states, 71, 87–88, 219–220, 221–222
- Vector addition, 26
- Vectors
 - basis (*see* Basis vectors)
 - column, 27–28, 29, 47, 49
 - concept of, 24–25
 - functions as, 238–245
 - normalized, 32, 40
 - orthogonal, 32, 64–67, 70
 - polarization, 91
 - quantum states and, 35–37
 - row, 29–30
 - three-(3-vector), 25, 27, 32–33, 74–75, 83
 - unit, 32, 40
 - See also* Bras (bra vectors); Eigenvectors; Kets (ket vectors)

- Vector space, 24–34
 - axioms, 24–27
 - bras, 28–30
 - column vectors, 27–28
 - functions and, 27–28, 239–240
 - inner products, 30–32
 - kets, 28–30
 - orthonormal bases, 32–34
 - tensor product as, 165
 - triangle inequality and, 142–146
- Velocity
 - momentum and, 286–288, 293
 - of quantum mechanical particle, 286–288
- Venn diagram, 14, 16
- Wave functions, 134–135, 236–238
 - action of Hamiltonian on, 320–321
 - calculating density matrices and, 206–207
 - collapse of, 126–127
 - entanglement and, 212–213
 - ground-state, 324–327
 - locality and, 225–226
 - measurement and collapsing, 218, 222–223
 - momentum and, 255–259
 - momentum representation, 260–265
 - of near-singlet state, 234
 - operator method and, 337–342
 - position representation, 254, 260–262, 263–265
 - of product state, 232
 - representing particles, 253–254
 - of singlet state, 233
 - state-vectors and, 136
- Wavelength, momentum and, 259–260
- Wave packets, 295–301
 - bimodal, 296–297
 - Gaussian, 301
 - harmonic oscillation and, 322
 - minimum-uncertainty, 301
 - moving at fixed speed, 276–277
 - for nonrelativistic free particle, 283
- Waves, 235–236
 - harmonic oscillator and, 313
- Wheeler, John, 52
- x axis, spins along, 41–42
- y axis, spins along, 42–45
- Zaxxon, 278
- Zero function, 239
- Zero operator, 133