FROM ETERNITY TO HERE

The Quest for the Ultimate Theory of Time

SEAN CARROLL

ONEWORLD
OXFORD

A Oneworld Book

First published in Great Britain and the Commonwealth
by Oneworld Publications 2011

A CIP record for this title is available
from the British Library

ISBN: 978-1-85168-795-4 (Trade Paperback)
ISBN: 978-1-85168-842-5 (Travel Edition)

Typeset by Jayvee, Trivandrum, India
Cover design by Richard Green
Printed and bound by

Oneworld Publications
185 Banbury Road
Oxford, OX2 7AR
England

CONTENTS

1

THE PAST IS PRESENT MEMORY

What is time? If no one asks me, I know. If I wish to explain it to one that asketh, I know not.

—St. Augustine, *Confessions*

The next time you find yourself in a bar, or on an airplane, or standing in line at the Department of Motor Vehicles, you can pass the time by asking the strangers around you how they would define the word *time*. That's what I started doing, anyway, as part of my research for this book. You'll probably hear interesting answers: "Time is what moves us along through life," "Time is what separates the past from the future," "Time is part of the universe," and more along those lines. My favorite was "Time is how we know when things happen."

All of these concepts capture some part of the truth. We might struggle to put the meaning of "time" into words, but like St. Augustine we nevertheless manage to deal with time pretty effectively in our everyday lives. Most people know how to read a clock, how to estimate the time it will take to drive to work or make a cup of coffee, and how to manage to meet their friends for dinner at roughly the appointed hour. Even if we can't easily articulate what exactly it is we mean by "time," its basic workings make sense at an intuitive level.

Like a Supreme Court justice confronted with obscenity, we know time when we see it, and for most purposes that's good enough. But certain aspects of time remain deeply mysterious. Do we really know what the word means?

WHAT WE MEAN BY *TIME*

The world does not present us with abstract concepts wrapped up with pretty bows, which we then must work to understand and reconcile with other concepts. Rather, the world presents us with phenomena, things that we observe and make note of, from which we must then work to derive concepts that help us understand how those phenomena relate to the rest of our experience. For subtle concepts such as entropy, this is pretty clear. You don't walk down the street and bump into some entropy; you have to observe a variety of phenomena in nature and discern a pattern that is best thought of in terms of a new concept you label "entropy." Armed with this helpful new concept, you observe even more phenomena, and you are inspired to refine and improve upon your original notion of what entropy really is.

For an idea as primitive and indispensable as "time," the fact that we invent the concept rather than having it handed to us by the universe is less obvious—time is something we literally don't know how to live without. Nevertheless, part of the task of science (and philosophy) is to take our intuitive notion of a basic concept such as "time" and turn it into something rigorous. What we find along the way is that we haven't been using this word in a single unambiguous fashion; it has a few different meanings, each of which merits its own careful elucidation.

Time comes in three different aspects, all of which are going to be important to us.

1. **Time labels moments in the universe.**
   Time is a coordinate; it helps us locate things.
2. **Time measures the duration elapsed between events.**
   Time is what clocks measure.
3. **Time is a medium through which we move.**
   Time is the agent of change. We move through it, or—equivalently—time flows past us, from the past, through the present, toward the future.

At first glance, these all sound somewhat similar. Time labels moments, it measures duration, and it moves from past to future—sure, nothing controversial about any of that. But as we dig more deeply, we'll see how these ideas don't *need* to be related to one another—they represent logically independent concepts that happen to be tightly intertwined in our actual world. Why that is so? The answer matters more than scientists have tended to think.

1. Time labels moments in the universe

John Archibald Wheeler, an influential American physicist who coined the term *black hole*, was once asked how he would define "time." After thinking for a while, he came up with this: "Time is Nature's way of keeping everything from happening at once."

There is a lot of truth there, and more than a little wisdom. When we ordinarily think about the world, not as scientists or philosophers but as people getting through life, we tend to identify "the world" as a collection of *things*, located in various *places*. Physicists combine all of the places together and label the whole collection "space," and they have different ways of thinking about the kinds of things that exist in space—atoms, elementary particles, quantum fields, depending on the context. But the underlying idea is the same. You're sitting in a room, there are various pieces of furniture, some books, perhaps food or other people, certainly some air molecules—the collection of all those things, everywhere from nearby to the far reaches of intergalactic space, is "the world."

And the world changes. We find objects in some particular arrangement, and we also find them in some other arrangement. (It's very hard to craft a sensible sentence along those lines without referring to the concept of time.) But we don't see the different configurations "simultaneously," or "at once." We see one configuration—here you are on the sofa, and the cat is in your lap—and then we see another configuration—the cat has jumped off your lap, annoyed at the lack of attention while you are engrossed in your book. So the world appears to us again and again, in various configurations, but these configurations are somehow distinct. Happily, we can label the various configurations to keep straight which is which—Miss Kitty is walking away "now"; she was on your lap "then." That label is time.

So the world exists, and what is more, the world *happens*, again and again. In that sense, the world is like the different frames of a film reel—a film whose camera view includes the entire universe. (There are also, as far as we can tell, an infinite number of frames, infinitesimally separated.) But of course, a film is much more than a pile of individual frames. Those frames better be in the right order, which is crucial for making sense of the movie. Time is the same way. We can say much more than "that happened," and "that also happened," and "that happened, too." We can say

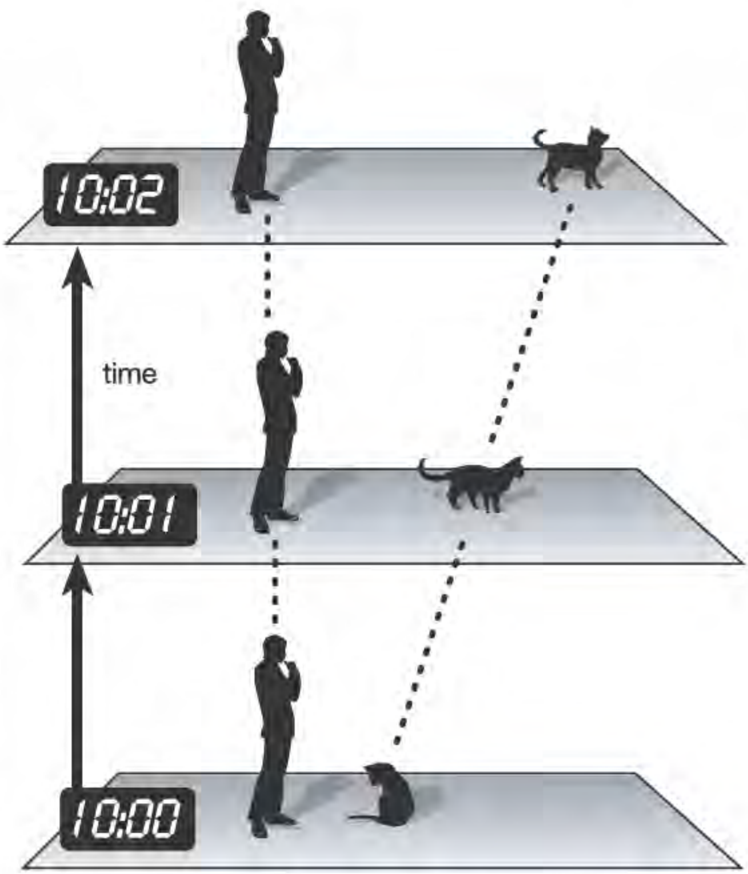

Figure 1: The world, ordered into different moments of time. Objects (including people and cats) persist from moment to moment, defining world lines that stretch through time.

that this happened *before* that happened, and the other thing is going to happen *after*. Time isn't just a label on each instance of the world; it provides a sequence that puts the different instances in order.

A real film, of course, doesn't include the entire universe within its field of view. Because of that, movie editing typically involves "cuts"—abrupt jumps from one scene or camera angle to another. Imagine a movie in which every single transition between two frames was a cut to a completely different scene. When shown through a projector, it would be incomprehensible—on the screen it would look like random static. Presumably there is some French avant-garde film that has already used this technique.

The real universe is not an avant-garde film. We experience a degree of continuity through time—if the cat is on your lap now, there might be some danger that she will stalk off, but there is little worry that she will simply dematerialize into nothingness one moment later. This continuity is not absolute, at the microscopic level; particles can appear and disappear, or at least transform under the right conditions into different kinds of particles. But there is not a wholesale rearrangement of reality from moment to moment.

This phenomenon of persistence allows us to think about "the world" in a different way. Instead of a collection of things distributed through space that keep changing into different configurations, we can think of the entire *history* of the world, or any particular thing in it, in one fell swoop. Rather than thinking of Miss Kitty as a particular arrangement of cells and fluids, we can think of her entire life stretching through time, from birth to death. The history of an object (a cat, a planet, an electron) through time defines its *world line*—the trajectory the object takes through space as time passes.[3] The world line of an object is just the complete set of positions the object has in the world, labeled by the particular time it was in each position.

Finding ourselves

Thinking of the entire history of the universe all at once, rather than thinking of the universe as a set of things that are constantly moving around, is the first step toward thinking of time as "kind of like space," which we will examine further in the chapters to come. We use both time and space to help us pinpoint things that happen in the universe. When you want to

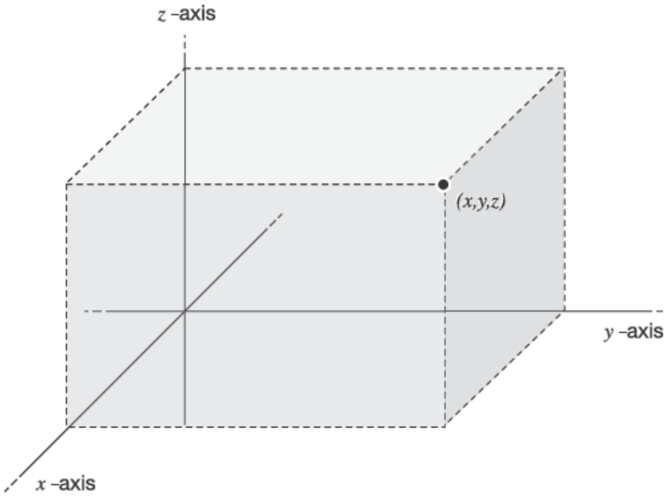

Figure 2: Coordinates attached to each point in space.

meet someone for coffee, or see a certain showing of a movie, or show up for work along with everyone else, you need to specify a time: "Let's meet at the coffee shop at 6:00 P.M. this Thursday."

If you want to meet someone, of course, it's not sufficient just to specify a time; you also need to specify a place. (Which coffee shop are we talking about here?) Physicists say that space is "three-dimensional." What that means is that we require three numbers to uniquely pick out a particular location. If the location is near the Earth, a physicist might give the latitude, longitude, and height above ground. If the location is somewhere far away, astronomically speaking, we might give its direction in the sky (two numbers, analogous to latitude and longitude), plus the distance from Earth. It doesn't matter how we choose to specify those three numbers; the crucial point is that you will always need exactly three. Those three numbers are the *coordinates* of that location in space. We can think of a little label attached to each point, telling us precisely what the coordinates of that point are.

In everyday life, we can often shortcut the need to specify all three coordinates of space. If you say "the coffee shop at Eighth and Main Street," you're implicitly giving two coordinates—"Eighth" and "Main Street"—and you're assuming that we all agree the coffee shop is likely to be at ground level, rather than in the air or underground. That's a convenience granted to us by the fact that much of the space we use to locate things in our daily lives is effectively two-dimensional, confined near the surface of the Earth. But in principle, all three coordinates are needed to specify a point in space.

Each point in space occurs once at each moment of time. If we specify a certain location in space at one definite moment in time, physicists call that an *event*. (This is not meant to imply that it's an especially exciting event; any random point in empty space at any particular moment of time would qualify, so long as it's uniquely specified.) What we call the "universe" is just the set of all events—every point in space, at every moment of time. So we need four numbers—three coordinates of space, and one of time—to uniquely pick out an event. That's why we say that the universe is four-dimensional. This is such a useful concept that we will often treat the whole collection, every point in space at every moment of time, as a single entity called *spacetime*.

This is a big conceptual leap, so it's worth pausing to take it in. It's natural to think of the world as a three-dimensional conglomeration that keeps changing ("happening over and over again, slightly differently each time"). We're now suggesting that we can think of the whole shebang, the entire history of the world, as a single four-dimensional thing, where the additional dimension is time. In this sense, time serves to slice up the four-dimensional universe into copies of space at each moment in time—the whole universe at 10:00 A.M. on January 20, 2010; the whole universe at 10:01 A.M. on January 20, 2010; and so on. There are an infinite number of such slices, together making up the universe.

2. Time measures the duration elapsed between events

The second aspect of time is that it measures the duration elapsed between events. That sounds pretty similar to the "labels moments in the universe" aspect already discussed, but there is a difference. Time not only labels and orders different moments; it also measures the distance between them.

When taking up the mantle of philosopher or scientist and trying to make sense of a subtle concept, it's helpful to look at things operationally—how do we actually use this idea in our experience?

When we use time, we refer to the measurements that we get by reading clocks. If you watch a TV show that is supposed to last one hour, the reading on your clock at the end of the show will be one hour later than what it read when the show began. That's what it *means* to say that one hour elapsed during the broadcast of that show: Your clock read an hour later when it ended than when it began.

But what makes a good clock? The primary criterion is that it should be consistent—it wouldn't do any good to have a clock that ticked really fast sometimes and really slowly at others. Fast or slow compared to what? The answer is: other clocks. As a matter of empirical fact (rather than logical necessity), there are some objects in the universe that are consistently periodic—they do the same thing over and over again, and when we put them next to one another we find them repeating in predictable patterns.

Think of planets in the Solar System. The Earth orbits around the Sun, returning to the same position relative to the distant stars once every year. By itself, that's not so meaningful—it's just the definition of a "year." But Mars, as it turns out, returns to the same position once every 1.88 years. That kind of statement is extremely meaningful—without recourse to the concept of a "year," we can say that Earth moves around the Sun 1.88 times every time Mars orbits just once.[4] Likewise, Venus moves around the Sun 1.63 times every time Earth orbits just once.

The key to measuring time is *synchronized repetition*—a wide variety of processes occur over and over again, and the number of times that one process repeats itself while another process returns to its original state is reliably predictable. The Earth spins on its axis, and it's going to do so 365.25 times every time the Earth moves around the Sun. The tiny crystal in a quartz watch vibrates 2,831,155,200 times every time the Earth spins on its axis. (That's 32,768 vibrations per second, 3,600 seconds in an hour, 24 hours in a day.[5]) The reason why quartz watches are reliable is that quartz crystal has extremely regular vibrations; even as the temperature or pressure changes, the crystal will vibrate the same number of times for every one rotation of the Earth.

So when we say that something is a good clock, we mean that it repeats itself in a predictable way relative to other good clocks. It is a fact about the universe that such clocks exist, and thank goodness. In particular, at the microscopic level where all that matters are the rules of quantum mechanics and the properties (masses, electric charges) of individual elementary particles, we find atoms and molecules that vibrate with absolutely predictable frequencies, forming a widespread array of excellent clocks marching in cheerful synchrony. A universe without good clocks—in which no processes repeated themselves a predictable number of times relative to other repeating processes—would be a scary universe indeed.[6]

Still, good clocks are not easy to come by. Traditional methods of timekeeping often referred to celestial objects—the positions of the Sun or stars in the sky—because things down here on Earth tend to be messy and unpredictable. In 1581, a young Galileo Galilei reportedly made a breakthrough discovery while he sat bored during a church service in Pisa. The chandelier overhead would swing gently back and forth, but it seemed to move more quickly when it was swinging widely (after a gust of wind, for example) and more slowly when wasn't moving as far. Intrigued, Galileo decided to measure how much time it took for each swing, using the only approximately periodic event to which he had ready access: the beating of his own pulse. He found something interesting: The number of heartbeats between swings of the chandelier was roughly the same, regardless of whether the swings were wide or narrow. The size of the oscillations—how far the pendulum swung back and forth—didn't affect the frequency of those oscillations. That's not unique to chandeliers in Pisan churches; it's a robust property of the kind of pendulum physicists call a "simple harmonic oscillator." And that's why pendulums form the centerpiece of grandfather clocks and other timekeeping devices: Their oscillations are extremely reliable. The craft of clock making involves the search for ever-more-reliable forms of oscillations, from vibrations in quartz to atomic resonances.

Our interest here is not really in the intricacies of clock construction, but in the meaning of time. We live in a world that contains all sorts of periodic processes, which repeat a predictable number of times in comparison to certain other periodic processes. And that's how we measure duration: by the number of repetitions of such a process. When we say that our TV program lasts one hour, we mean that the quartz crystal in our watch will oscillate 117,964,800 times between the start and end of the show (32,768 oscillations per second, 3,600 seconds in an hour).

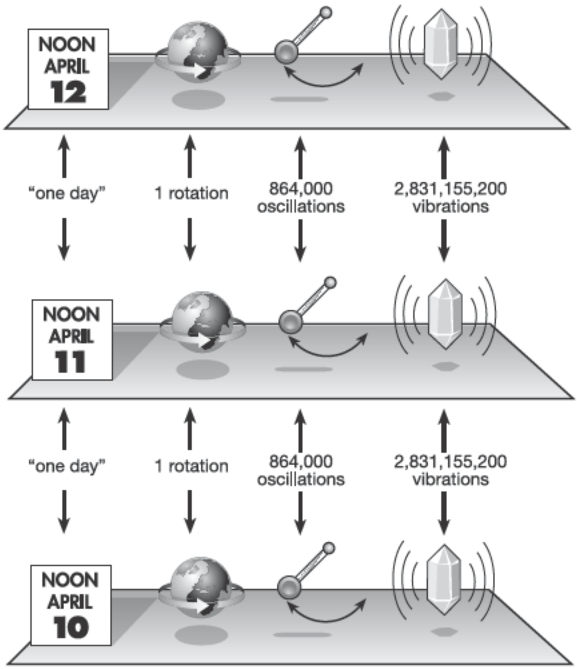

Figure 3: Good clocks exhibit synchronized repetition. Every time one day passes, the Earth rotates once about its axis, a pendulum with a period of 1 second oscillates 86,400 times, and a quartz watch crystal vibrates 2,831,155,200 times.

Notice that, by being careful about defining time, we seem to have eradicated the concept entirely. That's just what any decent definition should do—you don't want to define something in terms of itself. The passage of time can be completely recast in terms of certain things happening together, in synchrony. "The program lasts one hour" is equivalent to "there will be 117,964,800 oscillations of the quartz crystal in my watch between the beginning and end of the program" (give or take a few commercials). If you really wanted to, you could reinvent the entire superstructure of physics in a way that completely eliminated the concept of "time," by replacing it with elaborate specifications of how certain things happen in coincidence with certain other things.[7] But why would we want to? Thinking in terms of time is convenient, and more than that, it reflects a simple underlying order in the way the universe works.

Slowing, stopping, bending time

Armed with this finely honed understanding of what we mean by the passage of time, at least one big question can be answered: What would happen if time were to slow down throughout the universe? The answer is: That's not a sensible question to ask. Slow down relative to what? If time is what clocks measure, and every clock were to "slow down" by the same amount, it would have absolutely no effect at all. Telling time is about synchronized repetition, and as long as the rate of one oscillation is the same relative to some other oscillation, all is well.

As human beings we *feel* the passage of time. That's because there are periodic processes occurring within our own metabolism—breaths, heartbeats, electrical pulses, digestion, rhythms of the central nervous system. We are a complicated, interconnected collection of clocks. Our internal rhythms are not as reliable as a pendulum or a quartz crystal; they can be affected by external conditions or our emotional states, leading to the impression that time is passing more quickly or more slowly. But the truly reliable clocks ticking away inside our bodies—vibrating molecules, individual chemical reactions—aren't moving any faster or slower than usual.[8]

What could happen, on the other hand, is that certain physical processes that we thought were "good clocks" would somehow go out of synchronization—one clock slows down, or speeds up, compared to all the rest. A sensible response in that case would be to blame that particular clock, rather than casting aspersions on time itself. But if we stretch a bit, we can imagine a particular collection of clocks (including molecular vibrations and other periodic processes) that all change in concert with one another, but apart from the rest of the world. Then we would have to ask whether it was appropriate to say that the rate at which time passes had really changed within that collection.

Consider an extreme example. Nicholson Baker's novel *The Fermata* tells the story of a man, Arno Strine, with the ability to "stop time." (Mostly he uses this miraculous power to go around undressing women.) It wouldn't mean anything if time stopped everywhere; the point is that Arno keeps moving through time, while everything around him stops. We all know this is unrealistic, but it's instructive to reflect upon the way in which it flouts the laws of physics. What this approach to stopping time entails is that every kind of motion and rhythm in Arno's body continues as usual, while every kind of motion and rhythm in the outside world freezes absolutely still. Of course we have to imagine that time continues for all of the air and fluids within Arno, otherwise he would instantly die. But if the air in the rest of the room has truly stopped experiencing time, each molecule must remain suspended precisely in its location; consequently, Arno would be unable to move, trapped in a prison of rigidly stationary air molecules. Okay, let's be generous and assume that time would proceed normally for any air molecules that came sufficiently close to Arno's skin. (The book alludes to something of the sort.) But everything outside, by assumption, is not changing in any way. In particular, no sound or light would be able to travel to him from the outside world; Arno would be completely deaf and blind. It turns out not to be a promising environment for a Peeping Tom.[9]

What if, despite all the physical and narrative obstacles, something like this really could happen? Even if we can't stop time around us, presumably we can imagine speeding up the motion of some local clocks. If we truly measure time by synchronized repetition, and we arranged an ensemble of clocks that were all running fast compared to the outside world while they remained in synchrony with one another, wouldn't that be something like "time running faster" within that arrangement?

It depends. We've wandered far afield from what might actually happen in the real world, so let's establish some rules. We're fortunate enough to live in a universe that features very reliable clocks. Without such clocks, we can't use time to measure the duration between events. In the world of *The Fermata*, we could say that time slowed down for the universe outside Arno Strine—or, equivalently and perhaps more usefully, that time for him sped up, while the rest of the world went on as usual. But just as well, we could say that "time" was completely unaffected, and what changed were the laws of particle physics (masses, charges on different particles) within Arno's sphere of influence. Concepts like "time" are not handed to us unambiguously by the outside world but are invented by human beings trying to make sense of the universe. If the universe were very different, we might have to make sense of it in a different way.

Meanwhile, there is a very real way for one collection of clocks to measure time differently than another: have them move along different paths through space-time. That's completely compatible with our claim that "good clocks" should measure time in the same way, because we can't readily compare clocks unless they're next to one another in space. The total amount of time elapsed on two different trajectories can be different without leading to any inconsistencies. But it does lead to something important—the theory of relativity.

Twisty paths through spacetime

Through the miracle of synchronized repetition, time doesn't simply put different moments in the history of the universe into order; it also tells us "how far apart" they are (in time). We can say more than "1776 happened before 2010"; we can say "1776 happened 234 years before 2010."

I should emphasize a crucial distinction between "dividing the universe into different moments" and "measuring the elapsed time between events," a distinction that will become enormously important when we get to relativity. Let's imagine you are an ambitious temporal[10] engineer, and you're not satisfied to just have your wristwatch keep accurate time; you want to be able to know what time it is at every other event in spacetime as well. You might be tempted to wonder: Couldn't we (hypothetically) construct a time coordinate all throughout the universe, just by building an infinite number of clocks, synchronizing them to the same time, and scattering them throughout space? Then, wherever we went in spacetime, there would be a clock sitting at each point telling us what time it was, once and for all.

The real world, as we will see, doesn't let us construct an absolute universal time coordinate. For a long time people thought it did, under no less an authority than that of Sir Isaac Newton. In Newton's view of the universe, there was one particular right way to slice up the universe into slices of "space at a particular moment of time." And we could indeed, at least in a thought-experiment kind of way, send clocks all throughout the universe to set up a time coordinate that would uniquely specify when a certain event was taking place.

But in 1905, along comes Einstein with his special theory of relativity.[11] The central conceptual breakthrough of special relativity is that our two aspects of time, "time labels different moments" and "time is what clocks measure," are *not* equivalent, or even interchangeable. In particular, the scheme of setting up a time coordinate by sending clocks throughout the universe *would not work*: two clocks, leaving the same event and arriving at the same event but taking different paths to get there, will generally experience different durations along the journey, slipping out of synchronization. That's not because we haven't been careful enough to pick "good clocks," as defined above. It's because *the duration elapsed along two trajectories connecting two events in spacetime need not be the same.*

This idea isn't surprising, once we start thinking that "time is kind of like space." Consider an analogous statement, but for space instead of time: The distance traveled along two paths connecting two points in space need not be the same. Doesn't sound so surprising at all, does it? Of course we can connect two points in space by paths with different lengths; one could be straight and one could be curved, and we would always find that the distance along the curved path was greater. But the difference in *coordinates* between the same two points is always the same, regardless of how we get from one point to another. That's because, to drive home the obvious, the distance you travel is not the same as your change in coordinates. Consider a player in American football who zips back and forth while evading tacklers, and ends up advancing from the 30-yard line to the 80-yard line. (It should really be "the opponent's 20-yard line," but the point is clearer this way.) The change in coordinates is 50 yards, no matter how long or short was the total distance he ran.

The centerpiece of special relativity is the realization that *time is like that*. Our second definition, time is duration as measured by clocks, is analogous to the total length of a path through space; the clock itself is analogous to an odometer or some other instrument that measures the

Figure 4: Yard lines serve as coordinates on an American football field. A running back who advances the ball from the 30-yard line to the 80-yard line has changed coordinates by 50 yards, even though the distance traveled may have been much greater.

total distance traveled. This definition is simply not the same as the concept of a coordinate labeling different slices of spacetime (analogous to the yard lines on an American football field). And this is not some kind of technical problem that we can "fix" by building better clocks or making better choices about how we travel through spacetime; it's a feature of how the universe works, and we need to learn to live with it.

As fascinating and profound as it is that time works in many ways similar to space, it will come as no surprise that there are crucial differences as well. Two of them are central elements of the theory of relativity. First, while there are three dimensions of space, there is only one of time; that brute fact has important consequences for physics, as you might guess. And second, while a straight line between two points in space describes the shortest distance, a straight trajectory between two events in spacetime describes the *longest* elapsed duration.

But the most obvious, blatant, unmistakable difference between time and space is that time has a direction, and space doesn't. Time points from the past toward the future, while (out there in space, far away from local disturbances like the Earth) all directions of space are created equal. We can invert directions in space without doing damage to how physics works, but all sorts of real processes can happen in one direction of time but not the other. It's to this crucial difference that we now turn.

3. Time is a medium through which we move

The sociology experiment suggested at the beginning of this chapter, in which you ask strangers how they would define "time," also serves as a useful tool for distinguishing physicists from non-physicists. Nine times out of ten, a physicist will say something related to one of the first two notions above—time is a coordinate, or time is a measure of duration. Equally often, a non-physicist will say something related to the third aspect we mentioned—time is something that flows from past to future. Time whooshes along, from "back then" to "now" and on toward "later."

Or, conversely, someone might say that we move through time, as if time were a substance through which we could pass. In the Afterword to his classic *Zen and the Art of Motorcycle Maintenance*, Robert Pirsig relates a particular twist on this metaphor. The ancient Greeks, according to Pirsig, "saw the future as something that came upon them from behind their backs, with the past receding away before their eyes."[12] When you think about it, that seems a bit more honest than the conventional view that we march toward the future and away from the past. We know something about the past, from experience, while the future is more conjectural.

Common to these perspectives is the idea that time is a *thing*, and it's a thing that can *change*—flow around us, or pass by as we move through it. But conceptualizing time as some sort of substance with its own dynamics, perhaps even the ability to change at different rates depending on circumstances, raises one crucially important question.

What in the world is that supposed to *mean*?

Consider something that actually does flow, such as a river. We can think about the river from a passive or an active perspective: Either we are standing still as the water rushes by, or perhaps we are on a boat moving along with the river as the banks on either side move relative to us.

The river flows, no doubt about that. And what that means is that the location of some particular drop of river water *changes with time*—here it is at some moment, there it is just a bit later. And we can talk sensibly about the *rate* at which the river flows, which is just the velocity of the water—in other words, the distance that the water travels in a given amount of time. We could measure it in miles per hour, or meters per second, or whatever units of "distance traveled per interval of time" you prefer. The velocity may very well change from place to place or moment to moment—sometimes the river flows faster; sometimes it flows more slowly. When we are talking about the real flow of actual rivers, all this language makes perfect sense.

But when we examine carefully the notion that time itself somehow "flows," we hit a snag. The flow of the river was a change with time—but what is it supposed to mean to say that time changes with time? A literal flow is a change of location over time—but time doesn't have a "location." So what is it supposed to be changing with respect to?

Think of it this way: If time does flow, how would we describe its speed? It would have to be something like "*x* hours per hour"—an interval of time per unit time. And I can tell you what *x* is going to be—it's 1, all the time. The speed of time is 1 hour per hour, no matter what else might be going on in the universe.

The lesson to draw from all this is that it's not quite right to think of time as something that flows. It's a seductive metaphor, but not one that holds up under closer scrutiny. To extract ourselves from that way of thinking, it's helpful to stop picturing ourselves as positioned within the universe, with time flowing around us. Instead, let's think of the universe—all of the four-dimensional spacetime around us—as somehow a distinct entity, as if we were observing it from an external perspective. Only then can we appreciate time for what it truly is, rather than privileging our position right here in the middle of it.

The view from nowhen

We can't literally stand outside the universe. The universe is not some object that sits embedded in a larger space (as far as we know); it's the collection of everything that exists, space and time included. So we're not wondering what the universe would really look like from the point of view of someone outside it; no such being could possibly exist. Rather, we're trying to grasp the entirety of space and time as a single entity. Philosopher Huw Price calls this "the view from nowhen," a perspective separate from any particular moment in time.[13] We are all overly familiar with time, having dealt with it every day of our lives. But we can't help but situate ourselves within time, and it's useful to contemplate all of space and time in a single picture.

And what do we see, when looking down from nowhen? We don't see anything changing with time, because we are outside of time ourselves. Instead, we see all of history at once—past, present, and future. It's like thinking of space and time as a book, which we could in principle open to any passage, or even cut apart and spread out all the pages before us, rather than as a movie, where we are forced to watch events in sequence at specific times. We could also call this the Tralfamadorian perspective, after the aliens in Kurt Vonnegut's *Slaughterhouse-Five*. According to protagonist Billy Pilgrim,

> The Tralfamadorians can look at all the different moments just the way we can look at a stretch of the Rocky Mountains, for instance. They can see how permanent all the moments are, and they can look at any moment that interests them. It is just

> an illusion we have here on earth that one moment follows another like beads on a string, and that once a moment is gone it is gone forever.[14]

How do we reconstruct our conventional understanding of flowing time from this lofty timeless Tralfamadorian perch? What we see are correlated events, arranged in a sequence. There is a clock reading 6:45, and a person standing in their kitchen with a glass of water in one hand and an ice cube in the other. In another scene, the clock reads 6:46 and the person is again holding the glass of water, now with the ice cube inside. In yet another one, the clock reads 6:50 and the person is holding a slightly colder glass of water, now with the ice cube somewhat melted.

In the philosophical literature, this is sometimes called the "block time" or "block universe" perspective, thinking of all space and time as a single existing block of spacetime. For our present purposes, the important point is that we *can* think about time in this way. Rather than carrying a picture in the back of our minds in which time is a substance that flows around us or through which we move, we can think of an ordered sequence of correlated events, together constituting the entire universe. Time is then something we reconstruct from the correlations in these events. "This ice cube melted over the course of ten minutes" is equivalent to "the clock reads ten minutes later when the ice cube has melted than it does when the ice cube is put into the glass." We're not committing ourselves to some dramatic conceptual stance to the effect that it's *wrong* to think of ourselves as embedded within time; it just turns out to be more *useful*, when we get around to asking why time and the universe are the way they are, to be able to step outside and view the whole ball of wax from the perspective of nowhen.

Opinions differ, of course. The struggle to understand time is a puzzle of long standing, and what is "real" and what is "useful" have been very much up for debate. One of the most influential thinkers on the nature of time was St. Augustine, the fifth-century North African theologian and Father of the Church. Augustine is perhaps best known for developing the doctrine of original sin, but he was interdisciplinary enough to occasionally turn his hand to metaphysical issues. In Book XI of his *Confessions*, he discusses the nature of time.

> What is by now evident and clear is that neither future nor past exists, and it is inexact language to speak of three times—past, present, and future. Perhaps it would be exact to say: there are three times, a present of things past, a present of things present, a present of things to come. In the soul there are these three aspects of time, and I do not see them anywhere else. The present considering the past is memory, the present considering the present is immediate awareness, the present considering the future is expectation.[15]

Augustine doesn't like this block-universe business. He is what is known as a "presentist," someone who thinks that only the present moment is real—the past and future are things that we here in the present simply try to reconstruct, given the data and knowledge available to us. The viewpoint we've been describing, on the other hand, is (sensibly enough) known as "eternalism," which holds that past, present, and future are all equally real.[16]

Concerning the debate between eternalism and presentism, a typical physicist would say: "Who cares?" Perhaps surprisingly, physicists are not overly concerned with adjudicating which particular concepts are "real" or not. They care very much about how the real world works, but to them it's a matter of constructing comprehensive theoretical models and comparing them with empirical data. It's not the individual concepts characteristic of each model ("past," "future," "time") that matter; it's the structure as a whole. Indeed, it often turns out to be the case that one specific model can be described in two completely different ways, using an entirely different set of concepts.[17]

So, as scientists, our goal is to construct a model of reality that successfully accounts for all of these different notions of time—time is measured by clocks, time is a coordinate on spacetime, and our subjective feeling that time flows. The first two are actually very well understood in terms of Einstein's theory of relativity, as we will cover in Part Two of the book. But the third remains a bit mysterious. The reason why I am belaboring the notion of standing outside time to behold the entire universe as a single entity is because we need to distinguish the notion of time in and of itself from the perception of time as experienced from our parochial view within the present moment. The challenge before us is to reconcile these two perspectives.