This Boundary-less World

by Chris Fields

Separate individuals with separate identities through a shared linear time is the heart of classical physics that quantum theory rejects.

Ask yourself: can you determine, by observing the world around you, exactly which parts or aspects of the world are causing your current, right-now experience? Can you determine exactly which parts or aspects of the world caused any of your past experiences? It may take some reflection, but common sense will tell you that the answer to both of these questions is "no." Past experiences are, after all, in the past; what we observe now can tell us what *might have* caused a past experience, or what *probably* caused it, but only time travel could give us access to what actually did cause it. Is it enough, however, to make the past present? Does even right-now experience reveal its causes? Can you *demonstrate* that your current experience is being caused by the print on this page, and only by the print on this page? Can you prove to yourself that your experience is not being caused some other way, by some other means?

The question of whether the causes of experiences can be uniquely determined has traditionally been left to philosophers, or to fantasy writers like the authors of *The Matrix*. In everyday life, we assume that we know what is causing our experiences: if we see and hear someone speaking to us, for example, we assume that someone is really there, and is really speaking, and that their being there and speaking is what is causing our experience. Going about our daily business would be impossible without this assumption. Science clearly assumes this, too. Explaining the events that occur around us—the job of science—is largely a matter of explaining what caused them. The events we find most intriguing, or most troubling, are the ones for which we can find no satisfactory cause.

This chapter explores the consequences of questioning our everyday assumption that we know what the causes of our experiences are. It begins by examining the consequences of a second deep assumption: the assumption that, while how we describe the world matters to us, it does not matter to the world. This is an essential assumption of science; it is the assumption that there is a world “out there” that is independent of our human thought patterns and languages. While this is merely a practical, working assumption for most sciences, it is built into the very heart of quantum theory, the most accurate theory of the physical world that we have ever had. When combined with the assumption that we humans are *finite* observers, however, this deep assumption has a surprising consequence: it implies that we do not, and cannot, know the causes of our experiences precisely. It implies that we cannot know that two experiences have the same cause; hence we can never say, on the basis of experience, that we have encountered the same individual thing on two separate occasions. This insight allows us to understand, without having to do any mathematics, some strange predictions of quantum theory. It also allows us to understand why our brains devote so much energy to the task of *making up* connections between our memories and our present experiences. Understanding this, in turn, allows us to recognize time and individual identity for what they are: feelings we have about our experiences, feelings that help them make sense, but that hide their true character as unbounded, unpredictable processes. Embracing these unbounded and unpredictable processes leaves us with a sense of greater empathy, both for ourselves and for the world.

An old and deep assumption

What would the world be like if how we described it affected how it worked? It is perfectly ordinary to believe something rather like this, that how we describe the world affects how *we* work—how we act on and respond to the world. Naming a disease, for example, can dramatically change how we behave; that is why we value medical diagnoses. What would the world be like, though, if naming a disease changed how the *disease* behaved? This would be a strange world indeed!

One of the most basic assumptions of the sciences is that there is a way that the world works, whether we can figure it out or not. Our names for things, our ways of describing things, and our theories of how things work are important to us, but according to this basic assumption, they have no effect on the structure of the world or on how it works. Human languages and theories are just that: languages and theories. The world goes about its business completely independently of such things.

Turning this assumption into a scientific principle requires making it more precise. One way to make it more precise relies on the notion of “decomposing” the world into a collection of named, or even just noticed, things that are of interest for some purpose and on which attention is focused, together with “everything else” on which attention is not focused. Think, for example, of driving to work. One's attention is focused on a few things: the road, other cars, one's destination, perhaps the day's anticipated events. Most of the rest of the world remains unnoticed, unnamed, undifferentiated from the general background of “the world.” Such decompositions of “everything” or “the universe” into one or a small number of things—the “system(s) of interest” on which attention is focused together with a shared “environment” comprising everything else—is not just commonplace but essential in science. One can't measure everything that is going on, or from a broader perspective, any more than a tiny fraction of what is going on, so the “environment” that remains uncharacterized in any investigation is huge.

With this notion of decomposing the world into what is interesting for the present purposes and what is not, the basic assumption that there is a way the world works can be stated more precisely: how the world works is independent of our choices of decompositions. It is perhaps more forceful to state it as a symmetry principle that science must respect: all of the ways of decomposing the world that we do, can, or could employ are equivalent in at least one way: none of them affects how the world works. As we shall see, this principle of “decompositional equivalence” has some surprising consequences. Before examining these, however, let us look at some of the ways that this principle is used by the sciences, in particular by physics.

The basic idea behind physics, even in the ancient world, is that there are “laws of Nature” that apply in all situations, across the board. In the European Renaissance, when this kind of thinking could get you burned at the stake, the laws of Nature became “God's laws;” under that rubric, they still applied across the board and still could be figured out if one were clever enough. Laws that apply across the board are laws that do not depend on our names and descriptions. In a world in which the “laws of physics” were applied, in practice, mainly to planets and cannonballs, this subtle symmetry did not have much of an impact. However, as physics became more mathematical and experimental methods improved, this began to change. Newton's mathematical formulae applied to all things that had mass and were capable of motion; they said nothing about size or shape. Theorists quickly applied them to abstract entities, and experimentalists equipped with new methods tested the results. A classical understanding of fluid flow, for example, employs such an abstraction. One can imagine a volume element or “voxel” of fluid, embedded in the overall flow but enclosed within an imaginary boundary that prevents the exchange of material with the surroundings but allows the exchange of energy or momentum (Fig. 1). Deformations in the shape of this voxel correspond to changes in the pattern of the flow; these can be visualized and modeled with equations that predict measurable changes in pressure or flow velocity. No one would imagine, however, that drawing such a voxel boundary would change the physics of the flow; indeed, such imagined voxels are useful precisely because they can be drawn anywhere. Classical fluid dynamics, in other words, relies on and, hence, requires the assumption of decompositional equivalence.

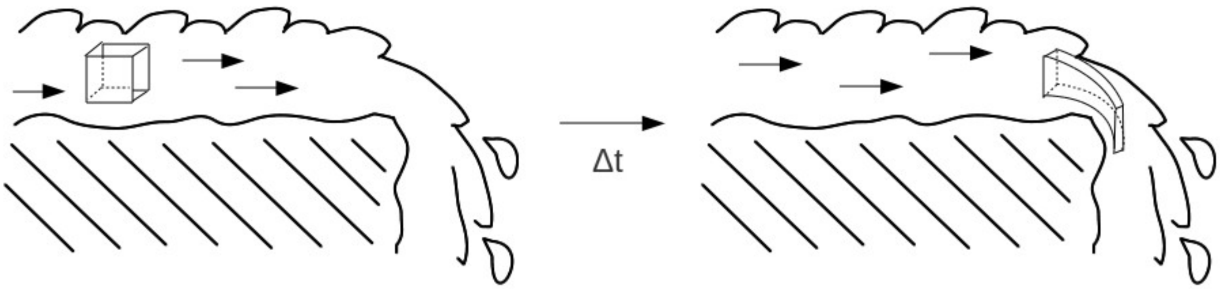

Fig. 1: A voxel of water flows across a waterfall during a time interval Δt.

> Decompositional equivalence lets us draw the voxel boundary *anywhere* without changing the physics of the flow.

The revival of the ancient Greek notion of “atoms” during the European Enlightenment further enhanced the centrality of decompositional equivalence within physics. Atoms had mass, and things made of atoms clearly could move; Newton's laws, therefore, applied to atoms. Suddenly, one could wonder: is the motion of an object not just the collective motion of its constituent atoms? Might the behaviors of planets and cannonballs be mathematical consequences, whether calculable in practice or not, of Newton's laws applied to their atoms? The new sciences of thermodynamics and chemistry seemed to confirm this: what the atoms in a thing were doing determined what the thing was doing. Could it be the case that the universe is just a bunch of atoms going about their business, and that everything we see—including ourselves—is just a description of this business from our idiosyncratic human perspective? Were this the case, decompositional equivalence would imply that our human perspective is not just idiosyncratic but one of many that are equally valid: that what we regard as “ordinary reality” is merely one of arbitrarily many ways of parsing the world. This idea that observed phenomena might be just surface features of an underlying, invisible dynamics—what philosophers would call “epiphenomenalism”—provided a third option in metaphysics, challenging both the “naïve” realist acceptance of ordinary objects at face value and the idealist rejection of the perceived world as an essentially illusory construction of the mind. It never reached a dominant status in philosophical circles, but it motivated scientists to ask whether Newton's laws were, in fact, obeyed by atoms.

The quantum revolution

Everyone now knows, of course, that Newton's laws are *not* obeyed by atoms. Phenomena at the atomic scale, the subatomic scale below it, and even the molecular scale above it are governed by the laws of quantum theory, the revolutionary new physics of Bohr, Einstein, Heisenberg and Schrödinger. In the world of quantum theory, electrons are both waves and particles, cats are both dead and alive, and only probabilities can be predicted. The world of quantum theory is, moreover, *our* world. The theory's predictions are regularly confirmed to accuracies of better than one part in 10 billion. Cell phones, digital cameras, and computers, to say nothing of nuclear weapons and power plants, all testify to the predictive success of quantum theory: all incorporate quantum technologies that depend on the violation of Newton's laws. Without the quantum revolution, the technological revolution of the twentieth century would not have happened.

What, however, *is* quantum theory? Textbooks present it as a morass of arcane mathematics. Both Bohr and Richard Feynman famously opined that anyone who was not bewildered by it did not understand it. The four basic principles of quantum theory can, however, be stated without recourse to complex mathematics. Two of these principles, as we will see, are consequences of the assumption of decompositional equivalence.

Understanding quantum theory requires understanding what is meant by the “physical state” of a system. The primary concern of physics is to explain, and if possible to predict, changes in this state, and the main difference between classical physics and quantum theory is in how they describe it. Suppose the system of interest is a chair. Classical physics describes the chair's physical state with two numbers, one representing its position in some coordinate system—its GPS coordinates, for example—and another representing its momentum. Momentum is a classical concept combining an object's speed, its direction of motion, and its mass; momentum is a measure of how much force an object can transfer to anything it collides with. Viewing the chair as a collection of atoms does not change the structure of this description: a position and a momentum are written down for each atom in the chair. The position of the chair is then the average position of its atoms, and the momentum of the chair is the sum of the momenta of its atoms. One can add additional numbers representing other “degrees of freedom” to the state description of the chair, for example, specifying its electrical charge or the speed with which it is rotating (its “angular momentum”), its temperature, the color it's painted, or any other observable property. One could also add these additional numbers to the state descriptions of its atoms. The list including all the numbers needed for the predictive task at hand is regarded as a “complete” description of the chair for the purposes of classical physics.

The construction of such a classical description of a physical state relies on an implicit assumption: the assumption that one can measure the position, momentum, electric charge, angular momentum, and any other physical properties of the chair that one wishes to write down. This implicit assumption, in turn, relies on two even deeper assumptions. One is that the chair can be identified independently of its position, momentum, etc. This is obviously required: if one can't *find* the chair, one can't measure its position. The other is that the properties to be measured can be measured independently. This is also obviously required. If measuring the position changed the momentum, for example, and measuring the momentum changed the position, one could never measure both and, hence, never write down a classical state description.

The development of quantum theory was motivated by many experiments, but one of the most basic motivations is that the second assumption required for the construction of classical state descriptions is violated at atomic scales. Measuring the position of an electron, for example, *does* change its momentum, and measuring its momentum *does* change its position: this is the content of Heisenberg's famous “uncertainty principle.” It is easy to see why this might be so: electrons are small, and measuring devices are big. Think of trying to measure the position of one of your hairs by having someone carefully approach it with a car. What Heisenberg's principle points out is that *any* measurement requires physical interaction, and any interaction that yields an answer in finite time involves an exchange of energy between the measurement device and the thing being measured. In the quantum world, such energy exchange is significant. Some pairs of measurements can be made in a way that the inevitable energy transfer does not pose a problem, but many cannot. Position and momentum, in particular, cannot be measured independently, so classical state descriptions are impossible.

Quantum state descriptions have a peculiar form that is dictated by the non-independence of measurements and the phenomenon of wave-particle duality. Instead of expressing the state of a system using just the last observed value of whatever properties are being measured, quantum theoretic state descriptions incorporate a representation of what the system could be doing during the period when it is not being observed. Why this is necessary is easiest to see by considering the phenomenon physicists call "spin," which is responsible for, among other things, the effectiveness of Polaroid® sunglasses. Measuring the spin—or "filtering" the polarization—of light in the vertical or "up" direction will create a beam of "up"-polarized light. Filtering the "up"-polarized light in the horizontal direction, at a right angle to the vertical, will yield no light at all. If, however, a filter oriented at an angle between vertical and horizontal is placed between the vertical and horizontal filters, some light will get through (Fig. 2). How could this be? The explanation is that the intermediate filter selects light with an intermediate polarization—one that incorporates "some" vertical and "some" horizontal light. Quantum theory makes this notion of "some" of the measurable alternatives being incorporated into a quantum state mathematically precise, and incorporates it into all quantum state descriptions. A quantum state describes, in other words, what is physically *possible*, and does so in a way that allows, given a mathematical specification of what will be measured, the calculation of what outcomes are probable.

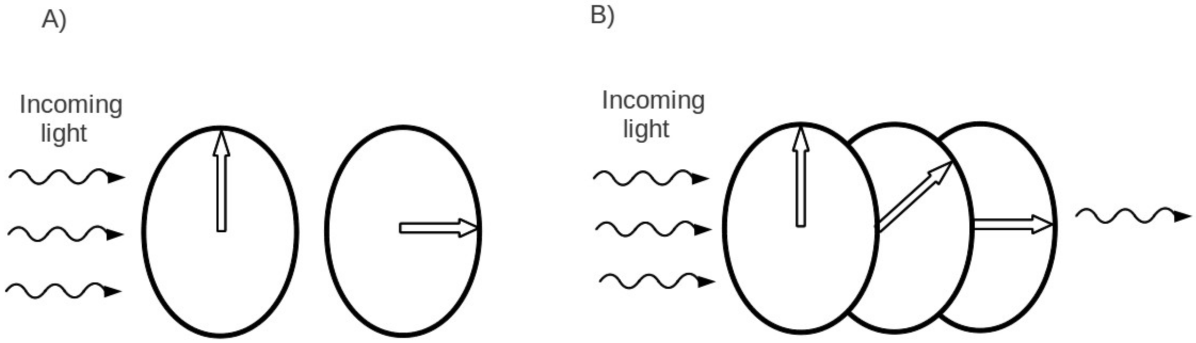

Fig. 2: Selecting first for vertical and then for horizontal spin lets no light through (Part A), but some light gets through if an intermediate selection is interposed between the two (Part B).

Given this concept of a quantum state, the basic assumptions of quantum theory can be stated as follows:

- Quantum states are descriptions of physical possibilities from which probabilities of observational outcomes can be calculated.
- Interactions between quantum states can be described mathematically as sums of the interactions between their components.
- Measurements of quantum states yield sets of real numbers, one for each independently measurable physical property.
- The quantum state of a composite system can be expressed mathematically as a product of the quantum states of its components.

That's it. From these assumptions, the *mathematics* of quantum theory follows. The physics of quantum theory is another matter. It is even more up for grabs now than it was when Bohr and Einstein had their famous debate at the 1927 Solvay Conference. Physical “interpretations” of the above assumptions range from the hyper-pragmatic “shut up and calculate” attitude of most academic physics classes to Henry Stapp's mind-matter dualism, Christoper Fuchs' notion that quantum theory is not physics at all but rather part of probability theory, the infinitely branching parallel worlds of Max Tegmark or David Wallace, or at least a half dozen other options. This question of physical interpretation can be set aside, however, in order to ask a simpler question: what does quantum theory tell us about observation? As we will see, what it tells us is that our earlier question about the causes of our experience is not merely metaphysical.

Heraclitus's revenge

The Greek philosopher Heraclitus (ca. 500 B.C.E.) is famous for saying “you cannot step into the same river twice” and for generally believing that our notion of identity over time—including our concept of our own identity over time—is deeply flawed. The empirical success of quantum theory demonstrates that Heraclitus was right. To see this, let us examine the second and fourth assumptions of quantum theory more closely. We will see that they are both statements of decompositional equivalence that have, in the context of the rest of quantum theory, some interesting consequences.

The second assumption of quantum theory is that interactions between components of a pair of quantum systems can be added together to calculate the interaction between the systems. This assumption seems straightforward: if I am interacting with my laptop, for example, it makes sense that the total interaction could be described by summing the interactions between my fingers and the keys, my eyes and the screen, and so on. Addition, however, is *associative*: it has the property that 2+(5+7) = (2+5)+7. One can insert pairs of parentheses into a sum any way one likes, as long as the grammar is preserved, and get the same answer. Applied to me and my laptop, this means that the components of me and the components of the laptop can be grouped any way one likes; one can group fingers and keys together, for example, and eyes and screen together, and calculate the interaction between the top half of the [me + laptop] system with the bottom half. The associativity of addition mathematically guarantees that this way of proceeding is not only just as good as any other way but that it gives exactly the same answer for the total interaction. The common-sense *boundary* between me and my laptop, in other words, makes *no difference at all to our physical interaction*; the boundary can be drawn in any of a nearly infinite variety of ways—think of all the ways to arrange our combined collection of atoms, for example—without changing a thing about the total interaction. This is decompositional equivalence with a vengeance! If the second assumption of quantum theory is true, then *none* of the boundaries we ordinarily regard the world as having makes *any difference at all* to physics. Thousands of extraordinarily detailed experiments tell us that this assumption is true.

The fourth assumption of quantum theory makes a similar point. Multiplication is also associative: otherwise one could not say that 2×9 and 3×6 are both equal to 18. Any way of defining the “components” of a quantum system that makes them multiply together to give the right system as their mathematical product is, therefore, the same as any other way as far as quantum theory is concerned. The universe, however, is by definition a quantum system: it is the system that quantum theory is about. So the universe can be divided up any way one pleases, as long as the “parts” so identified still fit together to form the universe. The particular boundaries that our ordinary perception tells us are so important, once again, make no difference at all.

This irrelevance of boundaries to physics tells us something striking. Ever since Plato, it has been a major goal of both philosophy and science to discover theories that “carve Nature at its joints” and reveal the “natural kinds” of objects that exist. If boundaries are irrelevant to physics, Nature *has no* physical joints, and, hence, there *are no* natural kinds as far as physics is concerned. What counts as a “kind of thing” is a matter of *our* classification scheme, not Nature's. This is sobering, and it becomes more sobering when one recognizes that our classification scheme is inevitably rough. The third assumption of quantum theory is that outcomes of measurements are real numbers. In practice, however, outcomes of measurements have to be recorded, and all recordings have finite size. So in practice, outcomes of measurements can be represented as finite *integers*, or as finite strings of 1s and 0s the way computers represent them. They can, moreover, be represented as *small* finite strings of 1s and 0s—strings 100 characters long, for example ($2^{100} \approx 10^{30}$). Our ability to distinguish the observable properties of systems is, therefore, much smaller than the extent to which such properties can vary as the boundaries around collections of atoms, for example, are subtly shifted. We cannot, in other words, distinguish between systems that are subtly different, or in many cases, between systems that are very different but nonetheless *look* similar. Our limited powers of observation mean that Nature can play tricks on us: if two systems that we cannot distinguish by observation were to be exchanged, we would be none the wiser. The world, therefore, has a symmetry: observable-dependent exchange symmetry. Given an observable—something we can measure—any number of distinct systems could be exchanged one for the other without us being able to tell the difference.

Observable-dependent exchange symmetry is Heraclitus's revenge. It tells us that even if we could step in the same river twice, we'd be unable to *demonstrate* that it was the same river. Put differently, just because it *looks like* the same river doesn't mean that it *is* the same river. A famous theorem of quantum theory—the Wooters-Zurek “no cloning” theorem—says that one cannot exactly duplicate a “pure” quantum state. Heraclitus's revenge tells us that even if we could duplicate a pure quantum state, we would not be able to demonstrate that we had. A copy that cannot, even in principle, be demonstrated to be a copy is, however, as good as no copy at all: the whole point of copying something is to *know* that one has made a copy. Another famous theorem of quantum theory—the Kochen-Specker “contextuality” theorem—tells us that the order in which we make measurements is significant, even for properties that can be measured independently. Heraclitus's revenge tells us that we can never be sure that we are measuring the same system, so the assumption that we would get the same results when making the measurements in different orders can never be relied upon. The most famous quantum theorem of all, Bell's theorem, faces physicists with a stark choice between accepting nonlocal causation—what Einstein called “spooky action at a distance”—and the curious claim that physical systems have no measurable properties except when they are actually being measured. Nonlocal causation is causation acting from somewhere else in the universe, somewhere that may be entirely unknown. Heraclitus's revenge tells us that we can never know whether a system we are examining is fully contained by any “local” boundary and, hence, never fully account for the causal influences acting on it. The assumption of locality—or in the language of physics, the assumption that a system can be “isolated from the environment” and hence isolated from potentially unknown causal influences—can, therefore, never be relied upon.

Quantum theory, in summary, elevates the assumption of decompositional equivalence from its classical status as an implicit methodological underpinning to the status of an axiom: not just one axiom, but two. It then relentlessly deduces the consequences of these axioms. What results is a description of the world in which observable-dependent exchange symmetry reigns supreme: a description that Heraclitus would have recognized, but that any metaphysics from Plato's onward finds strange and disconcerting. Experiments of ever-increasing complexity and sensitivity have told us, to date, that this description is accurate; indeed, the technological evidence of its accuracy surrounds us and drives our economy. We are, therefore, pressed to answer a further question: if the boundaries that we see make no difference to the physical processes by which the universe works, why do we see them? Even more critically, *how* do we see them? How can physical observers detect and come to rely upon boundaries that aren't really there?

Seeing boundaries

The question of why we see boundaries is, from a certain point of view, easy to answer: boundaries are inescapably invoked as soon as one says “we.” The very idea of a biological organism entails an organism-environment distinction and, hence, a boundary. Biological processes are processes that occur inside boundaries and that often have the preservation or maintenance of boundaries as their primary function. The idea that some organisms survive the accidents of life long enough to reproduce also invokes boundaries, both between those that survive and those that do not, and between parents and offspring. The process of cell division creates a boundary: one membrane-bounded cell becomes two, which then carry on as separate entities. Biology, in short, makes no sense without boundaries. It is not by accident that we view the boundaries formed by our respective skins as surpassingly important.

The question of how we see boundaries is more difficult; indeed, it has only been partially answered and has been answered at all only in the past few decades. Like most questions about how the brain works, the answer is complicated. It has, however, a simple summary. We see boundaries because we see *motion*. From the point of view of our brains, it is motion that creates objects. It is motion that creates identity through time, and, hence, motion that creates time itself.

From earliest infancy, human beings—indeed, most animals—are exquisitely sensitive to perceived motion. Think of walking down a path in the woods or through a secluded alleyway at night: any sudden motion will startle you, and will do so before you know what moved. Beginning in the 1970s, experimental studies began to show why. Our visual systems, which are very similar to those of other mammals, process perceived motion faster than they process perceived features like shape or color. We know something moved before we know what it is. From an evolutionary perspective, this makes sense. Snakes, thrown spears, and other dangerous things move fast, and it is better to avoid the occasional fast-moving safe thing than to fail to avoid the dangerous ones. As is usual in science, however, the question of *how* we accomplish this is both more interesting and harder to answer.

The “how?” question was first answered for monkeys, and the answer was quickly shown to be the same for humans. Our brains process information about motion and information about features—shape, color, etc.—using two different neural pathways (Fig. 3). The upper or “dorsal” visual pathway processes motion information: it determines where things are and how they are moving. It feeds this motion information directly to the motor system, which is why you jump first and ask questions later. The lower or “ventral” visual pathway processes feature information: it determines whether you are seeing a snake or a crooked branch, the face of a person or a face in the clouds. It feeds this feature information to the semantic memory system, which classifies what you are seeing into kinds of things. Motion information from the dorsal pathway is combined with—“bound to” in the language of neuroscience—feature information during this classification process. Classifying a moving object as a kind of thing is what puts a visual boundary around it; classifying something as an animal, for example, wraps it in an animal-like boundary. This classification and bounding process takes time. Classifying something as an animal as opposed to a non-animal takes about 200 milliseconds (ms), whereas noticing that it, whatever it is, is moving only takes about 50 ms. For comparison, becoming consciously aware of a scene containing a moving animal takes almost 250 ms. You jump, in other words, not just before you know what kind of thing is there, but before you are consciously aware that anything at all is there.

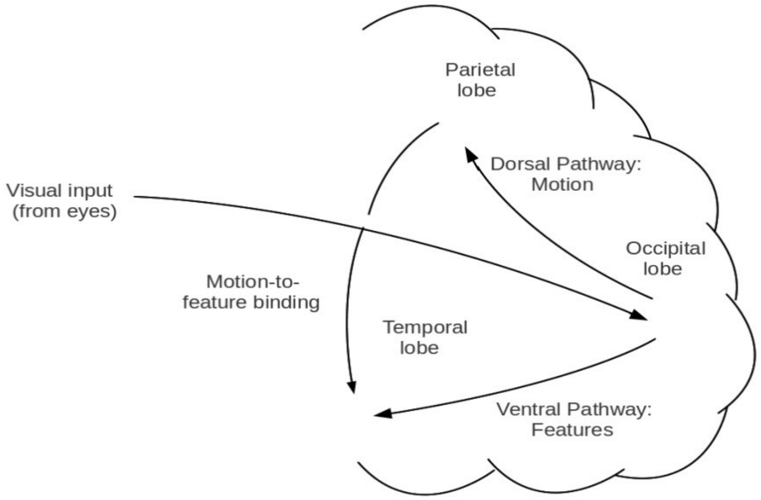

Fig. 3: Simplified cartoon of the back of the human brain, showing the processing of visual information from the primary visual cortex in the occipital lobe through the dorsal (motion) and ventral (feature) pathways. The pathways meet in the classification and memory-recording areas of the medial temporal lobe.

Every “how” answer, of course, leads to new “how” questions. How do we perceive motion, and how do we do it so fast? This time, the answer comes from studies that ask how—and under what conditions—we know whether something that has briefly disappeared is the same thing as something that has suddenly reappeared. Think, for example, of a bird that briefly disappears as it flies behind a tree. Such studies of “visual occlusion” show that people interpret some kinds of briefly-hidden motion as indicating that an object passed behind something and reappeared on the other side, while they interpret other kinds of briefly-hidden motion as indicating that one thing disappeared behind an occluding object and a *different* thing appeared on the other side (Fig. 4). Because the motion-analysis pathway works four times faster than the feature-analysis pathway, we make these snap judgments about sameness or difference—judgments about the *identity over time* of what we are seeing—independently of changes in the moving object's shape or color. Moreover, we make them without deliberation, well before we are consciously aware of what we are seeing.

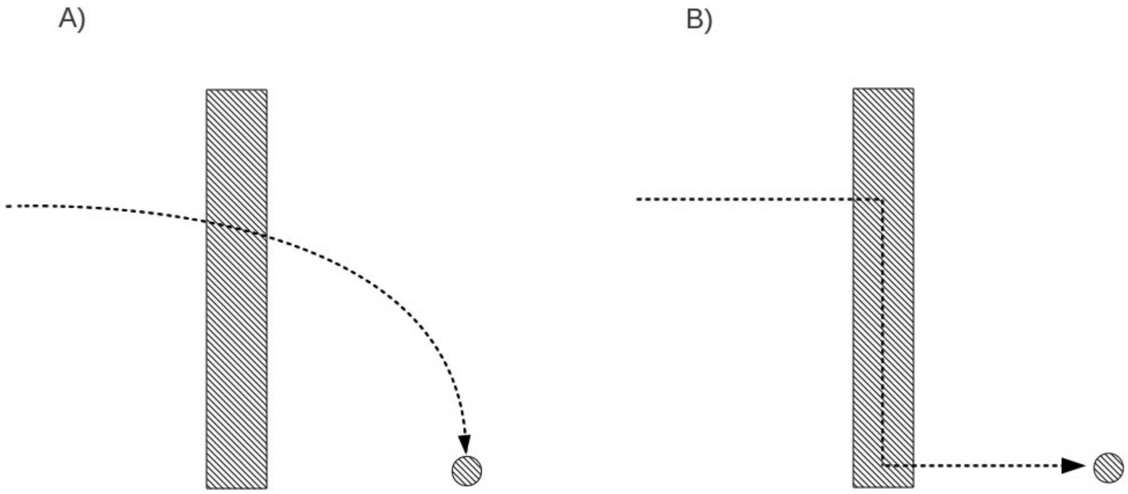

Fig. 4: Two trajectories that a moving object could follow. People see the smooth trajectory (A) as indicating object continuity; they see the angled trajectory (B) as indicating discontinuity, i.e., they see a “different” ball emerge from behind the occluder.

Creating identities

Studies of visual occlusion only tell us about identity judgments over short times—fractions of a second to seconds. Every day, however, we make identity judgments over much longer times. We wake up and recognize our spouse or partner, furniture, toothbrush, coffeepot, and countless other objects. We walk outside and recognize our house and car. We recognize relatives, colleagues, and buildings in distant cities after gaps in observation of months or years. We recognize them even after their appearance has changed in various ways. If you give some thought to our abilities to recognize objects of many kinds as being the very same things that were seen on previous occasions, it will seem inexplicable, indeed miraculous.

We can recognize things because we remember them. What does it mean, though, to remember a thing? We remember *events*—psychologists call this “episodic memory.” Events involve objects. How do we know that an object that was involved in some remembered event, for example, a colleague encountered at a meeting last year in Berlin, is the *very same thing* as an object observed today? What is the process that matches the memory, with its details of place, time, and context, with the current scene? We do not know, but evidence from both psychology and neuroscience suggests that identity over time is something that we *make up*. Specifically, we make up a plausible history that could explain how some single individual thing could have gotten from the event that we remember to the event that we are now experiencing. In cases where there *is no* plausible history—you encounter in New York a person who looks just like a friend you were just talking to, by telephone to Paris—you conclude that the object you are seeing is similar to the one you remember but not identical.

Making up histories is the job of a part of the brain—the posterior parietal cortex—that is also responsible for planning actions. We make up a history whenever we say how something happened, or why someone did something, without actually observing how it happened or asking why they did it. Making up histories is what allows us to explain things. It has, however, an even more crucial function. Our ability to make up histories is what allows us to recognize identity over time. In particular, it is what allows us to recognize *our own* identity over time, to connect our memories of *ourselves* on previous occasions to our experiences of ourselves right now. Without a sense of identity over time, moreover, we have little sense of time itself. Many of us know people whose memories are failing, and whose sense of the passage of time is slipping away as well.

If you find yourself suddenly worrying about Heraclitus's revenge, you are following the argument. In a world characterized by observable-dependent exchange symmetry, we can never know whether the histories we've made up are accurate or even close; we can never know whether the objects that we think are the very same things that we saw previously actually are. Indeed, from a purely physical point of view, we can be fairly sure that they aren't: every breath, for example, exchanges atoms in your body for atoms that used to be outside, and every meal does the same. What, then, is identity? Is it just something we make up? Are the things we see, with their histories and their tidy boundaries, merely fictions, just artifacts of our neurophysiology, our name-rich language, our time- and identity-based human culture?

Living in a boundary-less world

Physics and neuroscience both tell us that the answers to these existential questions is "yes." Physics and neuroscience both, of course, also assume as a practical matter that things and histories exist: otherwise experiments would be impossible. This contradiction between what we have to believe to do science and what the science we do tells us is true—a "dialetheia" in philosophical language—is something we just have to live with. It is part of living in a boundary-less world.

Accepting that our world is boundary-less changes very little but also changes everything. Our physics still works; indeed, it was our physics that forced us to accept boundary-lessness in the first place. Accepting that the world is boundary-less is nothing like classical idealism, the view that everything we see is just a product of our minds, a dream or an illusion. The boundary-less world is very much there: it is not a world without *stuff*, it is just a world without *boundaries*. If we smack it, it is hard—the electromagnetic forces that we experience as "hardness" are not illusions—and if it smacks us, we may very well be dead. It is perhaps better to say that a boundary-less world is a world in which boundaries are *irrelevant*. They can be drawn anywhere, for any reason, without making any difference at all to what the world is doing. The boundary-less world is a world in which all individual bounded things—including all of us—are imaginary constructs, fictions that may be useful or otherwise, created by ourselves, by others, or by some interaction or agreement between the two. It is a world, moreover, in which "imagination," "interaction," and "agreement" are just words for what the world is up to. It is a world in which what we, within our imagined boundaries, imagine ourselves to be doing is in fact what the *world* is doing, and what the world would be doing if all of our imagined boundaries were suddenly erased.

If it is the *world* that is doing, if it is the *world* that is observing, imagining boundaries, and agreeing, with itself, to talk in this or that way about "objects" or "individuals," then we have no choice but to regard the world as having the psychological properties that we regard ourselves as having. If we are agents, the world must be an agent. If we have free will, the world must have free will. If we are conscious, aware, having experiences, we must regard the world as conscious, aware, and having experiences. This panpsychist position—the view that the world itself, all of it, is conscious—is bracingly different from what most of us are used to. We humans tend to regard ourselves, our families, and our tribes as special. We regard our DNA, our cells, our environment as special. Our consciousness is the epitome of that specialness. The history of science, however, is a history of the deconstruction of our specialness, of our removal from the center of the universe. Quantum theory completes this process. In a boundary-less world, there is no specialness. There is only what is happening, what the world itself is doing. In such a world, panpsychism is not just respectful, it is the only self-consistent way to approach psychology.

What does it mean to live in a boundary-less world? I like to think of it with a *mantra*: “I am an experience that the universe is having.” This seems right: we are all just fleeting experiences. Hopefully we are pleasant ones, and useful.

Further reading

There are many excellent books on quantum theory by physicists. Lisa Randall's delightful *Warped Passages* includes a discussion of the history of quantum theory and is a good place to start. The four assumptions of quantum theory discussed in “The quantum revolution” are informal versions of the four axioms presented in Ch. 2 of *Quantum Computation and Quantum Information* by Nielsen and Chaung (Cambridge University Press, 2000), a standard textbook in the field. The sections on quantum theory and its various interpretations in Wikipedia and in the Stanford Encyclopedia of Philosophy are also good background reading and have many pointers into the relevant popular and professional literature: I suggest starting with http://en.wikipedia.org/wiki/Introduction_to_quantum_mechanics or http://plato.stanford.edu/entries/qt-idind/.

The discussion of “Heraclitus's revenge” above is presented more formally in my paper

“Bell's theorem from Moore's theorem,” *International Journal of General Systems* 42 (2013) 376-385. The discussion of “Seeing boundaries” is based on “Trajectory recognition as the basis for object individuation: A functional model of object file instantiation and object token encoding,” *Frontiers in Psychology* 2 (2011) 49 [12 pp]; that of “Creating identity” is based on “The very same thing: Extending the object token concept to incorporate causal constraints on individual identity,” *Advances in Cognitive Psychology* 8 (2012) 234-247. The inference from boundary-lessness to panpsychism is discussed more deeply in “A whole box of Pandoras: Systems, boundaries and free will in quantum theory,” *Journal of Experimental and Theoretical Artificial Intelligence* (in press, 2013). All of these have many pointers into the literature, and all are available for download at http://chrisfieldsresearch.com/.