Prologue

One summer many years ago, I was having a rough Sunday morning. It came after a string of bad days when I struggled with work and relationships, and I was quiet and tense. My eight-year-old son was invited to a birthday party about an hour's drive away, and as we left the house, I snatched a slim hardcover book from a pile near the door. Zachary was tired, and he napped on the way, quietly snoring in the back seat, and this was fine by me. When we got to the party, held in the backyard of a grand house in suburban Connecticut, I made small talk with the adults and then slipped away and sat under a tree and took out the book and started to read.

It was *The Origin of the Universe*, written by John Barrow, a theoretical physicist.[1] It began by describing Edwin Hubble's discovery that the universe was expanding, and then went over the evidence for the "Big Bang" theory of how everything started.

As I read, my heart began to beat faster. It was so exciting that we could know about all this, that I could be reading about events that happened fourteen billion years ago. Perhaps it's what people of faith feel like when reading Scripture—the experience of great truths being revealed. Learning about the universe, I felt insignificant, tiny in space and time. But I also felt proud of our species—that we could know so much about the incredibly long ago and incredibly far away, that we could make real progress on the most fundamental of all questions. And when the birthday party was over and I got up to get my son, the world was full of light.

Driving back, I talked to Zachary about what I learned, and as we spoke, I played with the fantasy of quitting my job as professor of psychology, getting a new degree, and becoming a cosmologist. But I was where I belonged. The tombstone of the philosopher Immanuel Kant has a quote from his *Critique of Pure Reason*: "Two things fill the mind with ever new and increasing admiration and awe, the more often and more steadily we reflect upon them: the starry heavens above me and the moral law within me." I had spent the morning being thrilled by the starry heavens; years later my research would turn to morality and moral psychology, and there again I would experience the same "admiration and awe" as Kant.

Honestly, though, just about all of psychology gives me this buzz. It's about the most interesting topic there is—*us*. It's about our feelings, experiences, plans, goals, fantasies, the most intimate aspects of our being.

The book you are holding is built from an Introduction to Psychology course that I've taught for many years as a professor at Yale University. This is one of the most popular courses at Yale, and I have taught thousands of undergraduates, sometimes as their very first course at a university. Based on these lectures, I created an online course that has had an enrollment, so far, of about a million students.[2]

I love teaching Introduction to Psychology. But there is a limit to how much one can convey in a series of lectures, and there is so much material to cover. And so I decided to write *The Human Mind*. The scope here is broad, and if you choose to read this from cover to cover, you'll have a grounding in every major aspect of the science of psychology. Among other things, *The Human Mind* will put forth the best answers we have to the following questions:

How does the brain—a three-pound lump of bloody meat—give rise to intelligence and conscious experience?
What did Freud get right about human nature?
What did Skinner get right about human nature?
Where does knowledge come from?
How does the mind of a child differ from that of an adult?
What is the relationship between language and thought?
How do our biases affect how we see and remember the world?
Are we rational beings?

What motivates us—and what is the purpose of feelings such as fear, disgust, and compassion?
How do we think of other people—including those from other social and ethnic groups?
How (and why) do we differ in personality, intelligence, and other traits?
What is the cause and treatment of different mental illnesses?
What makes people happy?

Each chapter of this book can be read as a stand-alone piece. It's fine if you decide to dive in and read about Freud, or language, or mental illness. Or even jump to the end, to the part on happiness. Nobody is judging you. But there is a flow to this book; there are themes and ideas that stretch across these disparate chapters, and there is a satisfaction to seeing the story unfold in its proper order.

Some parts of this story make people uncomfortable. We'll see that modern psychology accepts a mechanistic conception of mental life, one that is *materialist* (seeing the mind as a physical thing), *evolutionary* (seeing our psychologies as the product of biological evolution, shaped to a large extent by natural selection), and *causal* (seeing our thoughts and actions as the product of the forces of genes, culture, and individual experience).

You might worry that there is something missing here. This conception of mental life might seem to clash with commonsense notions of free choice and moral responsibility. It might seem to clash as well with the notion that humans have a transcendent or spiritual nature. The tension here is nicely illustrated by John Updike in his *Rabbit at Rest*, when Harry "Rabbit" Angstrom talks to his friend Charlie about Charlie's recent surgery:

> "Pig valves." Rabbit tries to hide his revulsion. "Was it terrible? They split your chest open and ran your blood through a machine?"
>
> "Piece of cake. You're knocked out cold. What's wrong with running your blood through a machine? What else you think you are, champ?"
>
> A God-made one-of-a-kind with an immortal soul breathed

in. A vehicle of grace. A battlefield of good and evil. An apprentice angel . . .

"You're just a soft machine," Charlie maintains.[3]

There are different ways to react to all this. I know philosophers and psychologists who confidently assert there's no such thing as free will or moral responsibility. And I've met others who reject the science, who worry that such an approach to the mind takes the specialness away from people, it diminishes us somehow. It's too reductionist, too crude. It reduces us to computers or lumps of cells or lab rats. They reason: "If psychology is going to tell me that I'm just a machine, that the most intimate aspects of my being are nothing more than neural firings, well, so much for psychology."

My own view is that we can find a middle ground here. I think the scientific perspective at the core of modern psychology is fully compatible with the existence of choice and morality and responsibility. Yes, we are, in the end, soft machines—but not *just* soft machines.

I want to end this prologue with a note of humility. We know so much about the physical world and so little about mental life. This isn't because physicists are smart and psychologists are stupid. It's because my chosen domain of study is so much harder than Barrow's. The mysteries of space and time turn out to be easier for our minds to grasp than those of consciousness and choice. In the pages that follow, I'll be honest about the limitations of our young science and critical of some colleagues who think we've solved it all.

But there's a real joy to being part of a young science. I find the study of psychology to be just as exhilarating as this study of the cosmos, and I hope you come to see it this way as well. We have made exciting progress in the field and I can't wait to talk about it. My fondest hope for this book is that the theories and discoveries reviewed here will give rise to a sort of awe in the reader, something akin to what I experienced when I read about the origins of the universe under that tree many years ago.

FOUNDATIONS

1

"Brain Makes Thought"

The Astonishing Hypothesis

In the late afternoon of September 13, 1848, something miraculous and terrible happened to the young foreman of a construction gang working in Cavendish, Vermont. Phineas Gage was preparing the roadbed for the laying of railway tracks, and he had a routine. He would bore holes into rocks, place explosive powder and a fuse inside the holes, and then pile sand and dirt over them. Then he would use a tamping iron—a piece of iron that looked like a javelin, about three and a half feet long and thirteen and a half pounds in weight—to pack it all down, making a plug over the explosive device. Later, the fuses would be lit, and the explosions would clear away the rocks.

Nobody knows what went wrong—perhaps something distracted him—but Gage slammed his tamping iron into a hole before he had poured the sand, and the blasting powder exploded. The iron shot upward with tremendous force. It entered the left side of Gage's jaw, passing behind the left eye through the left side of his brain and continuing out the top of his skull, landing yards away from him.

Gage lost consciousness—but just for a moment. His gang helped him onto an oxcart and took him to the Cavendish Inn, where he was renting a room. He sat on the veranda and told bystanders the story of what had just happened. When a medical expert finally arrived, Gage said, "Doctor, here is business enough for you."

It was touch and go for a while. Gage had an infection and required considerable treatment. But months later, he was seemingly recovered. He wasn't blind; he wasn't paralyzed; he retained the ability to speak and understand language; he didn't lose his intellectual capacities in any obvious way. You might think he was very lucky indeed.

But Gage wasn't lucky at all. As his doctor, John Martyn Harlow, wrote, Gage used to be: "the most efficient and capable man, a man of temperate habits, considerable energy of character, a sharp shrewd businessman." But afterward: "Gage was no longer Gage. He was fitful, irreverent, indulging at times in the grossest profanity manifesting but little deference for his fellows." He was, according to Harlow, "a child intellectually" with "the animal passions of a strong man."

Unable to return to his job as foreman, Gage held a string of jobs in the years that followed, including working as a stagecoach driver in Chile and being an attraction at Barnum's American Museum in New York, showing off his tamping iron and telling his story. Eleven years after his accident, he began to have seizures, and he died a few months later in his mother's home.

The story of Phineas Gage is a vivid illustration of how damage to the brain (and more specifically, damage to the frontal lobe, the part right behind the forehead) can have a profound influence on some of the most important aspects of who we are—our inhibitions, how we treat others, our character.

There has long been controversy over the details of what happened to Gage, and his story has become more and more extravagant as time has gone by.[1] But the account above is as accurate as I could make it. And anyway, the world contains thousands of Gages. There are many unfortunate cases of brain damage profoundly changing a person's nature.

So here's another one, with a different outcome. This is the story of Greg F., described by the neuroscientist Oliver Sacks, in an article called "The Last Hippie."[2] As a teenager, Greg was restless and rebellious. He dropped out of school and became a Hare Krishna, moving to a temple in New Orleans. After he spent some time there, the spiritual leader grew impressed with Greg, calling him a Holy Man. Then Greg slowly began to go blind. This was seen not as something to be treated, but as a spiritual event.

> He was "an illuminate," they told him; it was the "inner light" growing. . . . And indeed, he seemed to be becoming more spiritual by the day—an amazing new serenity had taken hold of him. He no longer showed his previous impatience or appetites, and he was sometimes found in a sort of daze, with a strange (many said "transcendental") smile on his face. It is beatitude, said his swami: he is becoming a saint.

After four years, the temple permitted his parents to visit Greg, and when they did,

> they were filled with horror: their lean, hairy son had become fat and hairless; he wore a continual "stupid" smile on his face (this at least was his father's word for it); he kept bursting into bits of song and verse, and making "idiotic" comments, while showing little deep emotion of any kind ("like he was scooped out, hollow inside," his father said); he had lost interest in everything "current"; he was disoriented—and he was totally blind.

It turned out that Greg had a tumor in his brain the size of an orange. It destroyed most of the parts of the brain devoted to vision and extended as well into his frontal lobes (up front) and his temporal lobes (on the sides). The tumor was removed but the damage was irreversible. Greg F. was worse off than Phineas Gage. He was not only blind; he lost most of his memory of decades of his life and was unable to form new memories. He was docile and without feeling, unable to survive on his own.

Phineas Gage, Greg F., and so many others illustrate what the Nobel Prize–winning biologist Francis Crick calls the Astonishing Hypothesis:

> You, your joys and your sorrows, your memories and your ambitions, your sense of personal identity and free will, are in fact no more than the behavior of a vast assembly of nerve cells and their associated molecules.[3]

There's a shorter version of this idea in one of Charles Darwin's notebooks: "Brain makes thought."[4]

The philosophical term for the position of Darwin and Crick is "materialism" (there is another meaning of this word that has to do with money; ignore this). For the materialist, there is nothing but physical stuff. There are no immaterial souls.

This is an odd and unnatural view.[5] People are more attracted to the doctrine of "dualism," which is that the mind (or the soul) is a fundamentally different kind of thing than the body. We are not one; we are two—bodies and souls. This is an idea that's present in most religions and most philosophical systems (Plato, for instance, was very much a dualist), but the most thoughtful and articulate defender of dualism was the philosopher René Descartes. In his honor, the idea that minds and bodies are distinct is often described as "Cartesian dualism."

Written in the early 1600s, one of Descartes' arguments for dualism had to do with the limitations of physical things. It might surprise you to hear this, but Descartes was familiar with robots. He had visited the French Royal Gardens, which the philosopher Owen Flanagan describes as "a veritable seventeenth-century Disneyland,"[6] and was impressed by the automata driven by hydraulic force:

> You may have seen in the grottoes and fountains which are in our royal gardens that the simple force with which water moves in issuing from its source is sufficient to put into motion various machines

> and even to set various instruments playing or to make them pronounce words accordingly to the varied disposition of the tubes which convey the water. . . . In entering [strangers] necessarily tread on certain tiles or places, which are so disposed that if they approach a bathing Diana, they cause her to hide in the rosebuds, and if they try to follow her, they cause Neptune to come forward to meet them threatening them with his trident.[7]

There is an analogy here to the human body; the springs and motors in the robots correspond to muscles and tendons; the tubes in the robots correspond to nerves. So, one can wonder: Are we nothing more than complicated machines?

Descartes said no. This analogy, he argued, was correct for nonhuman animals. Their actions are solely the product of their physical constitution. They are nothing more than *bêtes machines*—beast machines. There are chilling stories, perhaps apocryphal, of Descartes participating in operations on live dogs—vivisection—presumably believing that their shrieks of agony were akin to the noises that broken machines sometimes make. Without souls, after all, they are incapable of feeling.

But humans are different. We see this in the unpredictable nature of our actions. The doctor taps your knee and you kick your lower leg, and, yes, this can be done by your body alone, under the same principles that govern the motions of robot Diana and robot Neptune. But you can also *choose* to kick out your leg right now, just for the hell of it. This is the sort of willful action that Descartes believed a physical thing could never do. And so he concluded: We are not physical things.

Descartes' other main argument for dualism is better known. He starts with the question "What can we know for sure?" and answers: Not much. You believe you were born in such and so place, but maybe you were lied to. Perhaps, as some children fantasize, you are of royal blood, and it's only due to some misadventure that you got stuck with the disappointing family of commoners who raised you. Or, to get really weird, maybe the universe was created five seconds ago, and your memories are all false. This is unlikely, but it's possible.

You might believe that you are now in a certain physical environment, sitting on a chair with your loyal hound by your side, my book in one hand and a cigar in another (or whatever). But Descartes observed that we often believe such things when we are dreaming. You might protest that you are not dreaming right now, but most dreamers don't know they are dreaming.

You can be wrong that you have a body. Philosophers have long worried that our experience might be an illusion created by the devil, and a modern version of this concern is played out in my favorite movie, *The Matrix*, which imagines a world in which everyday human experience is an illusion created by a malevolent computer. Some philosophers take this even further and argue that we are parts of computer simulations—essentially video-game characters. You might agree with me that this seems batty, but how can we know for sure?

We can't. But Descartes notes there's one thing that cannot be doubted—our own existence as thinking beings. The famous line is *Cogito ergo sum*—I think, therefore I am. You might not be sure you have a body, but you can be sure that there is a you who is asking the question. Drawing on this distinction between how we think about minds and how we think about bodies, Descartes concludes,

> I knew that I was a substance the whole essence or nature of which is to think, and that for its existence there is no need of any place, nor does it depend on any material thing. . . . That is to say, the soul by which I am what I am, is entirely distinct from body.[8]

This feels right. Our gut intuition is that we are not our bodies; we inhabit these bodies. We are *Ghosts in the Shell*, in the evocative phrase of manga artist Masamune Shirow. This is why we so easily create and understand fictions where bodies and souls come apart. Think about Franz Kafka's *Metamorphosis*, which begins with: "As Gregor Samsa awoke one morning from uneasy dreams he found himself transformed in his bed into a gigantic insect." Or the scene from *The Odyssey*, in which the goddess Circe transforms Odysseus' men into pigs: "They had the head, and voice, and bristles, and body of swine; but their minds remained unchanged as before. So they were penned there, weeping." Or countless other tales of possession, body swaps, frightening or lovable ghosts, and the like.[9]

Dualism has an appealing real-world consequence. If you are not your body, you can survive its destruction. Maybe you'll end up in some spirit world, or ascend to heaven, or occupy some other body. Now, there are clever ways in which materialists can also arrive at some sort of afterlife belief—perhaps God could somehow reanimate your corpse, repairing it as one would a broken watch. But for the most part materialism is a grim doctrine, tying your survival to the fate of your all-too-fragile flesh.

With all the arguments in favor of dualism, then, and all its attractions, why are modern-day psychologists so confident that its opposite—materialism—is correct?

Let's go back to Descartes' arguments. He was correct about the limitations of material things—hundreds of years ago. But now we have an expanded understanding of what such things are capable of. For Descartes, the idea that a machine can do something as complicated as playing chess would be ludicrous. This requires rational deliberation; it's not a matter of reflex. But of course, there are now machines that play chess, better than any human. One might reasonably wonder about other limits of physical things—can computers *feel*?—and we will get to these doubts later on. But the point here is that Descartes' argument no longer flies. The complexity of our actions is not proof of dualism.

As for what Descartes could and could not imagine, many philosophers have pointed out that he was too quick to assume that such a conceptual exercise could tell us about how things really are. Yes, you can doubt that you have a body and can imagine yourself without one. But this doesn't mean that this is possible. After all, I can imagine a spaceship moving faster than light—there are many in science fiction. Descartes' method reflects how we think about minds, not what's true about minds.

Consider all the problems with dualism and all the evidence against it. Talking about an immaterial Cartesian soul, the psychologist Steven Pinker writes, "How does the spook interact with solid matter? How does an ethereal nothing respond to flashes, pokes, and beeps and get arms and legs to move?"[10] This is an old complaint. In 1643, Elizabeth Stuart, the former queen of Bohemia, wrote to Descartes to complain how hard it is to take seriously the idea that "an immaterial thing could move and be moved by a body."[11]

To be fair, the other option, that brains make thoughts, can be equally hard to stomach. Here's Gottfried Leibniz in 1712: "In imagining that there is a machine whose structure would enable it to think, feel, and have perception, one could think of it as enlarged yet preserving its same proportions, so that one could enter into it as one does a mill. If we did this, we should find nothing within but parts which push upon each another, we should never see anything which would explain a perception."[12] More than a few modern neuroscientists have become tempted by dualism toward the end of their careers, often with a similar argument to that of Leibniz—they've spent their lives studying the brain, and they found no physical sign of consciousness residing there, and so perhaps it's in the spirit realm after all.

In the end, what decides the issue is all the evidence that the brain *is* implicated in thought, though not in a way that's apparent to someone peering at an opened skull. The change in Phineas Gage's character was caused by a very physical tamping iron going through his very physical head. And of course, you didn't have to wait until 1848 to appreciate that a blow to your skull can affect your consciousness and your memory, and can obliterate them permanently if the blow is hard enough. Everybody knew that dementia could rob you of your rationality or that coffee and alcohol can, in different ways, inflame the passions. (As Pinker puts it: "The supposedly immaterial soul, we now know, can be bisected with a knife, altered by chemicals, started or stopped by electricity, and extinguished by a sharp blow or by insufficient oxygen."[13])

What is new is that we can now observe the brain at work. You can put someone in a brain scanner, for instance, and tell from the parts of the brain that are active whether they are thinking about their favorite song or the layout of their apartment or a mathematical problem. We may not be too far from the point where we can look at the brain of a sleeping person and know what they are dreaming about.[14]

Is there any hope for Descartes' position? There are important distinctions between mental events and physical events, and some contemporary philosophers defend what they describe as mild forms of dualism.[15] These views are worth discussing, though not here. But almost nobody defends the sort of hard-core dualism maintained by Descartes, so-called substance dualism, where the mind is a different kind of stuff than the brain, where the process of thinking occurs in an immaterial realm, separate from the laws of nature. This theory is dead as a theory can be.

Sentient Meat

Okay, then, what is the physical seat of thought? What is the source of our emotions, decision making, our passions, our pains, and everything else? Even a dualist must address some version of the question; the soul must connect to some part of our physical being to make the body act and receive its sensory information. (Descartes believed the conduit to be the pineal gland.)

For most of history, people thought that the answer was the heart. This was apparently the belief of diverse populations around the world, including the Maya, the Aztecs, the Inuit, the Hopi, the Jews, the Egyptians, the Indians, and the Chinese. It is a view at the foundation of Western philosophy; Aristotle wrote, "And of course, the brain is not responsible for any of the sensations at all. The correct view [is] that the seat and source of sensation is the region of the heart. . . . The motions of pleasure and pain, and generally all sensation plainly have their source in the heart." After all, the heart responds to feelings; it pounds when you are lustful or angry; it is still when you are calm.[16]

But the brain is also a serious contender. There are psychological experiments suggesting that the commonsense view is that our consciousness is located above the neck.[17] In work that I've done with my colleague and wife, Christina Starmans, we find that even young children, when asked in various ways to locate "where" a person is, tend to answer that their real location isn't the chest; it's right between the eyes.[18]

Head or heart? Over history, this has been the focus of much debate, nicely captured by a line from *The Merchant of Venice*, written in the late 1500s:

> Tell me where is fancy bred,
> Or in the heart or in the head?[19]

As you probably have heard, we now know the answer—it's the head. The brain is a mere one fiftieth of our body weight but consumes about a quarter of the calories we burn off when we are at rest—it's an energy hog. The human brain is also humongous. Baby heads are bowling balls, which is one reason why human females, relative to females of other species, have such a prolonged and painful childbirth.

If you've never seen a brain, you might imagine that it would look impressive. It's often described as the most complicated thing in the known universe, after all. Perhaps it would glow, maybe there would be flashing colored lights or something. But no, it's just meat. One can eat brain—I've had it with cream sauce (not human brain, mind you—you shouldn't eat human brain; you can get this terrible disease, kuru, which is much like mad cow disease, and it's one reason not to be a cannibal). When you take it out of the head, the brain is dull gray; inside the head it is bright red because of all the blood.

There is a science fiction short story by Terry Bisson that nicely captures just how strange this is.[20] The story is in the form of a dialogue by a pair of hyperintelligent aliens traveling through the universe to find sentient beings, and then they find us:

> "Meat. They're made of meat."
>
> "Meat?"
>
> "There's no doubt about it. We picked up several from different parts of the planet, took them aboard our recon vessels, and probed them all the way through. They're completely meat."
>
> "That's impossible. What about the radio signals? The messages to the stars?"
>
> "They use the radio waves to talk, but the signals don't come from them. The signals come from machines."

"So who made the machines? That's who we want to contact."

"They made the machines. That's what I'm trying to tell you. Meat made the machines."

"That's ridiculous. How can meat make a machine? You're asking me to believe in sentient meat."

The aliens agree to erase the records and report that our solar system is unoccupied.

To explore the mystery of "sentient meat," we will start small, with neurons, and work our way up. By weight, most of the brain is fat and blood, and there are cells in the brain other than neurons (about half the brain is composed of glial cells, which support, clean up after, and nourish neurons). But the story of mental life is fundamentally the story of neurons, which is why the study of the biological basis of thought is called *neuroscience*.[21] Below are the parts of a neuron.

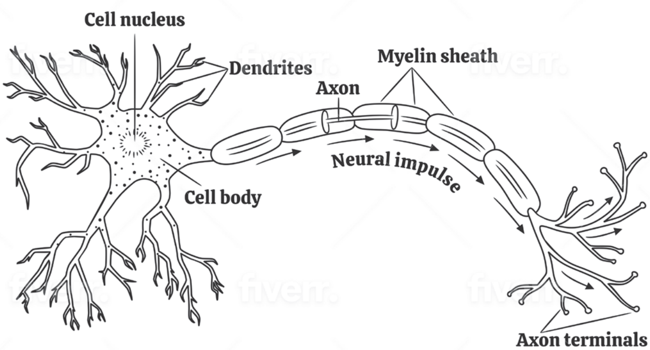

Like other cells, neurons have a cell body that keeps the cell alive and that houses a nucleus that contains the chromosomes that are made up of DNA. The cell body also coordinates the inputs from other neurons or from the senses. It gets this information through dendrites, which protrude from the cell body like tree branches—"dendrite" comes from the

Greek word meaning "tree." If the cell body receives the right sort of input from these dendrites, it causes the neuron to fire, and then an electrical signal goes down a long part of the neuron known as the axon. While dendrites are tiny, axons are long; there are single axons that run from your spinal cord all the way to your big toe. Axons have myelin sheaths—coatings of fatty tissue—that work like insulation on a wire, making the communication in the neuron run more efficiently. Diseases such as multiple sclerosis involve damage to the myelin sheath, which causes problems with action, perception, and thought.

The neuron then communicates with other neurons or, less frequently, to organs and muscles.

Summing up, then, here's how information typically flows:

dendrites > cell body > axon > the dendrites of other neurons

Some neurons are sensory neurons and take in information from the external world; some are motor neurons and go out to the external world. If you touch something hot, and you feel pain, that is because of sensory neurons: If you reach for something, that's due to the working of motor neurons. Other neurons—interneurons—don't connect directly to the world; they connect to one another, and this is how the thinking happens.

There's a puzzle here: When neurons talk to these other neurons or connect to the world, their communication is all or nothing. Neurons fire or they don't fire. It's like a gun—the bullet doesn't go faster if you pull the trigger with all your strength. But perception and action are graded. You can feel the difference between touching a warm plate and a hot stove; you can poke someone gently or really hard.

The solution to the puzzle is that assemblies of neurons have certain ways to represent the intensity of experience and action. One is the number of neurons that fire. If N neurons correspond to a mild experience, then N × 100 neurons may correspond to an intense experience. There is also the frequency of firing of individual neurons; an individual neuron might denote a mild sensation with fire . . . fire . . . fire . . . fire; and an intense sensation with firefirefirefirefirefire. Similar means of coding explain how motor neurons can code for intensity, allowing you to choose to pound the wall with your fist or gently stroke your newborn's cheek.

One major finding about neurons was discovered by the neuroscientist Santiago Ramón y Cajal in the 1800s. I've said that neurons talk to one another when the axon of one communicates with the dendrite of another. But neurons don't touch. There is a tiny gap between the axon of one neuron and the dendrite of another—typically about 20 to 40 nanometers. This gap is known as a *synapse*.

One of the great scientific disputes of the last century was over how the message got across this passage. This was known as the War of the Soups and the Sparks; the options were chemical (soup) or electrical (spark).[22] Long story short, the Soups won. As Cajal discovered, when neurons fire, axons release chemicals that we now call neurotransmitters; these cross the synapses to act on the dendrites of other neurons.

I said before that the cell body decides whether to fire based on the inputs of the dendrites impinging on it. Now I can spell this out a bit. The effect of these incoming neurotransmitters can either be excitatory, which is that they increase the likelihood of a neuron firing, or inhibitory, so that they bring down the likelihood of a neuron firing. The cell bodies put this all together, calculating whether all the increases and decreases add up to enough of a sum to fire.

Neurotransmitters are a large part of the story of how the brain works, and they have considerable practical import. We have invented drugs that interact in different ways with the workings of the neurotransmitters, and these can treat diseases, enhance pleasure, or increase focus.

Or kill. For an example of some deadly interaction, take curare, a drug used by some Indigenous people in South America while hunting—it's put on the tip of a dart or arrow. Curare is an antagonist, which means that it makes neurotransmitters less available for use. More specifically, it inhibits sensitivity to a neurotransmitter called acetylcholine, which is how motor neurons communicate with muscles. This is how curare paralyzes the prey. In large enough doses, it kills, because motor neurons also keep an animal breathing. Conveniently, curare is safe to eat, so you can dine on an animal that you felled with a curare-tipped dart.

Other drugs are agonists; they increase the availability of neurotransmitters for the brain to use. More specifically, they work on neurotransmitters such as norepinephrine that are involved with arousal, increasing euphoria, wakefulness, and control of attention. This is how (in different ways, and to different extents) drugs like speed and Ritalin and cocaine work.

This is what thinking is, then: neurons talking to other neurons via neurotransmitters. By one estimate, the brain of an adult human contains about eighty-six billion neurons, each connecting to thousands or tens of thousands of other neurons, leading to hundreds of trillions of connections, a combinatorial explosion that's mind boggling.[23]

But how does this give rise to experience? How does fire, fire, [not fire], fire, [not fire], etc., give rise to laughing at an excellent tweet or grieving the death of someone you love? And what about action? Our brains are physical things, but they are wired up so that they guide us to act in ways that seemingly transcend the laws of physics. William James puts it like this:

> If some iron filings be sprinkled on a table and a magnet brought near them, they will fly through the air for a certain distance and stick to its surface. . . . But let a card cover the poles of the magnet, and the filings will press forever against its surface without its ever occurring to them to pass around its sides. . . . If now we pass from such actions as these to those of living things, we notice a striking difference. Romeo wants Juliet as the filings want the magnet; and if no obstacles intervene he moves towards her by as straight a line as they. But Romeo and Juliet, if a wall be built between them, do not remain idiotically pressing their faces against its opposite sides like the magnet and the filings with the card. Romeo soon finds a circuitous way, by scaling the wall or otherwise, of touching Juliet's lips directly. With the filings the path is fixed; whether it reaches

> the end depends on accidents. With the lover it is the end which is fixed, the path may be modified indefinitely.[24]

Other creatures with brains have similar capacities for feelings and for rational action. A chimpanzee might shake with fear or bellow with rage. A cheetah chasing an antelope who darts behind a tree won't run into the tree but will move around it. How do brains do all this?

People often get caught up in the seemingly paradoxical nature of this inquiry—isn't it weird that we're using our brains to understand our brains? One physicist, Emerson M. Pugh, wrote, "If the human brain were so simple that we could understand it, we would be so simple that we couldn't." The comedian Emo Phillips says, "I used to think that the brain was the most wonderful organ in my body. Then I realized who was telling me this."

Charles Darwin added a twist to this, implying that the human brain was the *second* most interesting thing in nature. What could be more marvelous? Well, look at the ground below you:

> It is certain that there may be extraordinary mental activity with an extremely small absolute mass of nervous matter: thus the wonderfully diversified instincts, mental powers, and affections of ants are notorious, yet their cerebral ganglia are not so large as the quarter of a small pin's head. Under this point of view, the brain of an ant is one of the most marvellous atoms of matter in the world, perhaps more so than the brain of a man.[25]

There is reason for optimism, though, when it comes to understanding how brains make us—and ants—smart. When I talked about Descartes' dualism, I noted that computers prove that brute physical things are capable of capacities we associate with intelligence. One can take this further now and consider that computers work through simple processes that do stupid things like turn 0 to 1 or 1 to 0. If you have enough processes of this sort and they are put together in the right way, then chess playing, mathematical ability, language parsing, and all the rest arise. And here we meet up with neuroscience, as these binary operations look intriguingly like the basic dichotomy expressed by neurons inside the brain—fire versus not fire. This is progress, then: Computers suggest that the project of neuroscience is feasible, that intelligence can arise from the proper interaction of components that are themselves entirely unintelligent. As the polymath Alan Turing speculated in the 1940s, the human mind may be a computing machine.[26]

The mind as a computer? There is a dismissive reaction you sometimes get here: We used to see brains as hydraulic machines or clocks; then as telegraphic networks; then as telephone exchanges; and now, finally, we see them as computers. Perhaps this is yet another metaphor, a way of talking that will one day be supplanted by something else.

I agree that seeing the brain as akin to a Mac or a PC is just a metaphor, and not a very good one. Neurons are much slower to communicate than parts of computers are, and the brain is "wired up" differently than the computer I'm using to write these words. Much of the brain operates simultaneously—in parallel—whereas computers are largely serial.

And there are a thousand more specific ways in which computers function differently than brains. To take a case we'll explore later in our discussion of memory, when you ask someone a question about a past experience, the question itself can alter their recollection of the scene. If you show someone a movie and later ask, "Did you see the children getting on the school bus?" the person is more likely to remember, later on, a school bus in the scene, even if there wasn't one.[27] Indeed, repeated questioning can lead to the creation of false memories. Computers don't work this way. You can search for "school bus" a hundred times; this will not create a file with "school bus" on your hard drive. Human memory and computer memory work in very different ways.

But there is another sense in which the brain really is a computer. The brain processes information—it *computes*. Not long ago, at the time that Alan Turing did his pioneering work that led to the foundation of artificial intelligence, "computer" referred to a sort of person, someone who worked at the job of computing. To call the brain a computer in this sense isn't a metaphor. It is an interesting claim. It means that it carries out mathematical and logical calculations, it manipulates symbols. Calculating that one plus one is two is computation, and so is reasoning that if all men are mortal and Socrates is a man, then Socrates is mortal. The idea that brains are computers in this sense has shaped psychological theories of mental life, and we'll return to it when we talk about capacities such as language and perception.

Thinking about the brain as a computer has an interesting implication: Just as the study of computation can inform us about psychology, studies of the mind can help us build better computers. If you want to build machines that can walk in a straight line, recognize faces, and understand language, it's sensible enough to check out how people do it, in the same way that Leonardo da Vinci studied the wings of birds to figure out how to make a flying machine.

The brain is not just a large bowl of porridge. It contains parts that do different things. These parts sometimes get called areas, systems, modules, or faculties—the linguist Noam Chomsky called them "mental organs" to emphasize how they can be as different from one another as the organs below the neck, like the kidney or the spleen.[28]

Actually, the idea that mental life has parts has been popular since before anyone even knew about the brain. Plato, for instance, talked of a trinity—a "spirit" that lives in the chest and is involved in righteous anger, the "appetite" located in the stomach and related to desires, and "reason," in the head (at last!), which oversees the other two.

One attempt to subdivide the brain came from Franz Josef Gall, who founded the school of phrenology. Gall had a very good idea and a very bad idea. The good idea was that the different parts of the brain were specialized for different things, many of which would make a contemporary neuroscientist nod with agreement, such as number, time, and language. Gall's ideas were popular in the early 1800s and left us with these lovely diagrams in which the skull is depicted with dotted lines segmenting different parts, like the drawings of cows that one sometimes sees in steakhouses—marking off chuck, sirloin, round, and so on—except that here the parts are mental traits and capacities.

The bad idea—phrenology—was that these brain areas bloat up the more they are used, and this causes bumps on the skull. Someone skilled in the tools of phrenology, Gall claimed, can put their hands on people's heads and learn about their characters. Phrenology used to be quite the thing. Karl Marx was a convert, and he would sometimes rub the heads of people he met. Queen Victoria was similarly entranced and hired phrenologists to test out the skulls of her children.[29]

I don't need to tell you that this is a quite goofy notion. But in Gall's insistence that different parts of the brain have different functions, and not just general functions like reason or appetite, but specific ones, like language, he was a scientist ahead of his time.

If the brain is composed of parts, then we can learn about how it works by taking it apart. This idea was nicely expressed in 1669 by the anatomist Nicolaus Steno:

> The brain being indeed a machine, we must not hope to find its artifice through other ways than those which are used to find the artifice of the other machines. It thus remains to do what we would do for any other machine; I mean to dismantle it piece by piece and to consider what these can do separately and together.[30]

Neuroscience can be said to have properly begun when scholars started to put this strategy into practice, by looking at those sad cases where natural causes did the dismantling. In 1861, a French physician named Paul Broca discovered a patient who was intelligent and could fully understand what was said to him, but who could only produce one word, "tan," which he said no matter what was said to him, usually twice in a row—"tan, tan." After he died an autopsy found brain damage in part of his frontal lobe, now known as Broca's area.

Years later, the neurologist Carl Wernicke discovered a patient with a different language disorder—she had problems understanding speech and could talk rapidly and fluently, but what she said was gibberish. This was related to another part of the brain, located in the back of the temporal lobe, usually on the left side of the brain, which has come to be known as Wernicke's area. (Note that finding the precise locations of these areas is of practical value; when doctors are cutting into brains during surgery, they want to avoid hitting areas that serve valuable functions.)

Language provides one illustration of how different parts of the brain have different capacities. Let's look at others, taking a brief tour of the brain.

The cortex is the part on the surface, right under the skull, and many parts of the brain that are highly relevant to our mental lives are subcortical, meaning that they lie just below the cortex. A pretty metaphor I've heard envisions the brain as a peach—the skin is the cortex, and the subcortical structures are parts of the stone. (The flesh of the peach is the white matter, largely composed of glial cells.) Such subcortical structures include:

The medulla, which controls automatic functions like heart rate, blood pressure, and swallowing.

The cerebellum, which is involved in movement, posture, motor learning, and certain aspects of language. (To get a sense of how complex these systems are, note that the cerebellum contains about thirty billion neurons.)

The hypothalamus, which is involved in sleep and wakefulness and hunger and thirst and sex. (This corresponds best to what Plato was talking about when he speculated about the appetitive part of the soul, although he located it in the stomach.)

The limbic system, involved with the emotions.

The hippocampus, involved in long-term memory storage and memory of locations and things in space.

The pituitary gland, which secretes hormones involved in sex and reproduction and other things, and is interesting to historians of science and philosophy because, according to Descartes, it served as the conduit between body and soul.

Now let's turn to the skin of the peach. The first thing you notice when you look at a brain is that it's all wrinkly. This is because it's crumpled up. If you were to take a brain, pull out the cortex, and smooth it out, it's about two feet square.

The cortex breaks down into different lobes. You have the frontal lobe (conveniently enough on the front), the parietal lobe, the occipital lobe, and the temporal (next to the temple!) lobe.

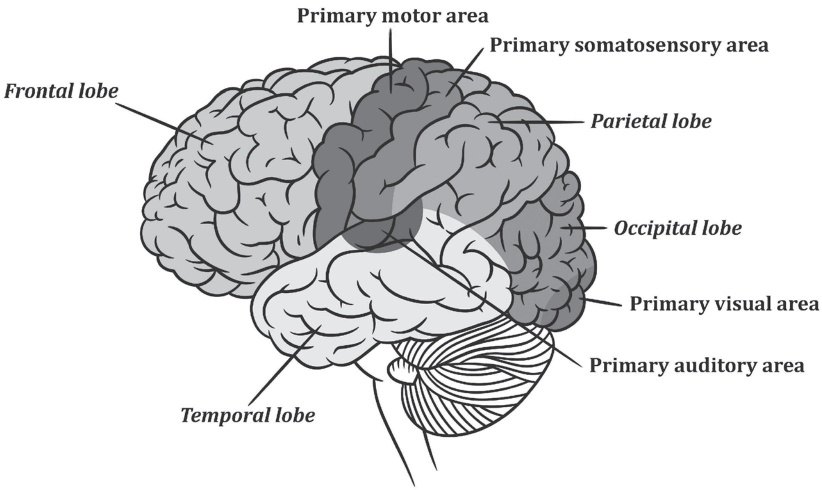

Some parts of the cortex contain "maps"—areas of the brain that correspond to parts of the body. If you give mild electrical shocks to neurons in the primary motor area, associated parts of the body twitch accordingly, while shocks to the primary somatosensory (soma = body) area lead to sensations in the corresponding areas. They are called maps because they are isomorphic with the body. For instance, the part of the brain representing the right index finger is close to the part of the brain representing the right thumb, which is close to the part of the brain representing the right wrist.

While the organization of these motor and sensory maps matches the organization of the body, the size does not—rather, brain size corresponds to the extent to which there is a lot of motor control or sensory discrimination. For instance, the part of the brain corresponding to the hand is bigger than the part corresponding to the chest because there is so much more sensation going on in your hand than in your chest, and so it gets more brain.

Aside from these maps, much of the rest of the cortex is involved in higher-order functions, such as language, reasoning, and moral judgment. Fish don't have any cerebral cortex, reptiles and birds have a little bit, mammals have more, and primates, including humans, have a lot.

How do we know which parts of the cortex do what? We've already mentioned studies where electrical impulses are applied to parts of the brain, but these are unusual, typically done with people who are having some sort of brain surgery. Far more common are those methods that look at the real-time brain activity of healthy people with intact heads. One popular technique is fMRI, which uses a strong magnetic field to look at the distribution of blood flow to the brain, seeing which parts are active when people think about different things. It is this that has the possibility of, in an almost literal sense, reading minds.

More techniques are emerging all the time. One method doesn't scan the brain at all—it *influences* it. This is TMS—transcranial magnetic stimulation, which uses magnetic fields to stimulate cells in the brain. Apply TMS to one part of the brain and it can impair language; apply it to other areas and it can cause the body to move without the person's volition. (I got to experience this once myself while I was visiting a laboratory in Kyoto—it felt weird to have my fingers twitch for no reason.)

We also know a lot about the brain from so-called natural experiments when people have tumors or strokes or accidents, as with unfortunate individuals such as Phineas Gage or Greg F. From these tragic cases, we can learn about which parts of the brain correspond to which functions, helping us to understand the correspondences between mind and brain.

As one example, some types of brain damage lead to agnosia—disorders of perception. Those with agnosia can see just fine, but often fail to recognize objects. When shown a picture, they can often describe the parts, but can't recognize how these parts make up the whole. More specifically, there is prosopagnosia, where one can't recognize faces. Oliver

Sacks wrote a classic book many years ago called *The Man Who Mistook His Wife for a Hat*.[31] This was a series of profiles of people who had surprising neurological disorders, including a man whose prosopagnosia was so bad that, as it says in the title, he couldn't distinguish his wife's face from a hat. In more common and milder forms, someone suffering from prosopagnosia can recognize faces as faces, but they can't recognize whose faces they are. This all illustrates the distinction between sensation and perception, something we'll get to later.

If you just look at a brain—if you remove it from somebody's head and put it on a table—it looks symmetrical. But it isn't. The asymmetry of the brain is manifested in handedness. Some people are right-handed, and others are left-handed, and since motor control comes from the brain, this suggests that the brain itself is asymmetrical. It has a right side and a left side, and they are not identical.

The difference between the sides is often exaggerated in popular articles. There is no such thing as "right-brain" people and "left-brain" people. Most of the functions of the brain are on both sides of the brain.

Still, there are differences. The left side of the brain is usually more involved with language and with the capacity for reason and logic, while the right side of the brain is more involved with social processes, imagination, and music. Some of these right-left differences are inborn; others are produced by experience. As one striking case of how culture shapes our brains, learning to read—a relatively recent human invention—reconfigures the brain, creating a region that is active when looking at words (called a "letterbox") in the left hemisphere and shifting the processing of faces more to the right hemisphere.[32]

The halves of the brain connect to the world in accord with a principle of contralateral organization, which means that due to a quirk of evolutionary history that's not entirely understood, your right brain sees the left side of the world and your left brain sees the right side of the world; your right hemisphere controls the left side of the body and your left hemisphere controls the right side of the body.

Suppose, then, that a psychologist presented a picture very quickly to your left visual field, too fast for your eyes to turn to see it head-on, and asked you to name the image. The information would go to the right side of your brain. Since typically the left side of the brain is involved in the processing of language, there would be a fraction of a second delay for you to name the picture, because the information has to make it to the left side, where the names are stored. If it was flashed to the right visual field, you would be just a little bit faster.

Information goes from one half of the brain to the other mostly through the corpus callosum—a network of neurons in the middle of your skull. If you think of the two halves of the brain as a city bisected by a river, like Budapest with the river Danube running down the middle, the corpus callosum is like thousands of small bridges uniting the city.

What if you were to cut the corpus callosum? This used to be done as a last resort for extreme cases of epilepsy. Epilepsy could be viewed as an electrical storm in the brain; the idea of this radical surgery was to isolate and shrink the electrical storms. And it really did help with the seizures. It also meant that the two halves of the brain couldn't readily communicate with each other, and this had some serious consequences.

In one case, for instance, a split-brain patient would find herself putting on clothes with her right hand and removing them with her left; in another, a patient who was shopping would put something in the shopping cart with one hand and take it out with the other. The left hand of another patient would suddenly strike his wife; another patient's left hand tried to choke him. This sort of behavior is so common in split-brain patients that it has a name—alien hand syndrome.[33] Separated, the two parts of the brain no longer act in unison; they can be seen as two individuals occupying, and sometimes fighting over, the same body.

Some scientists and philosophers draw a disturbing conclusion from the split-brain cases. They argue that for all of us, including those with intact corpus callosa, each half of our brain can be seen as a separate person. There is the language-using you, the one who is mostly in charge, and this is who is reading these words. But there is another you, a silent partner, also conscious, sitting next to the language-using self. When you split the brain, you liberate this silent self from the dominant "you," and the two selves may fight it out for control. But this radical conclusion is controversial, and there's no consensus as to what's really going on in the mind (or minds) of someone with a split brain.[34]

This brief tour of the brain has come to an end. We will go on to explore, throughout the rest of the book, psychological processes such as decision making, memory, and emotional experience. We know that all of these are the consequences of activity in brains—neurons working in concert with other neurons. This is an amazing discovery.

But sometimes people get *too* interested in the brain, neglecting the mind. Occasionally you bump into a neuroscientist who says that theirs is the real science. Sure, you can talk about ideas, emotions, short-term memory, and so on, but when you really get down to it, the serious theories will be about brain areas, neurons, and neurotransmitters. This is what matters. Neuroscience makes psychology irrelevant.

This attack on psychology is based on a confusion about how scientific explanation works. Just because we know about molecular biology doesn't mean that we stopped talking about hearts, kidneys, respiration, and the like. The sciences of anatomy and physiology did not disappear. Cars are made of atoms but understanding how a car works requires appealing to higher-level structures such as engines, transmissions, and brakes, which is why physics will never replace auto mechanics. Or to take an analogy closer to psychology, you can best understand the strategies that a computer uses to play chess by looking at the program it implements, not the material stuff the computer is made of. The same chess program can run on a 1980s mainframe computer, a 1990s desktop, a 2000 laptop, or a present-day smartphone. The physical structure of the hardware changes with each generation, but the program can stay constant.

If your neuroscientist is skeptical, ask how they would respond to a physicist telling them that the real science of the mind is ultimately about atoms and molecules, which are themselves composed of elementary particles—so why are they wasting their time talking about neurons and glial cells and the hippocampus, and the like? The neuroscientist would promptly protest that certain important scientific findings, such as the discovery that the hippocampus is involved in memory storage or that lack of dopamine is implicated in Parkinson's disease, can't be captured in the language of physics. And this would be a good response. Well, the psychologist can tell the neuroscientist much the same thing.

Indeed, it turns out that one can do psychology without studying the brain, even though the mind *is* the brain. Some of the most astonishing findings in our field have been done by scholars who couldn't tell a neuron from a nematode. And while we're at it, one can do psychology without studying evolution, even though the brain has evolved, and one can do psychology without studying child development, even though we were all once children. Many routes understanding has, as Yoda would put it.

I think part of the enthusiasm about brains reflects the commonsense dualism we talked about earlier. Occasionally when boasting about the effects of some sort of therapeutic or educational intervention, people claim, "It changes the brain!" But *everything* changes the brain. Reading this sentence just changed your brain, because you're thinking about it and thinking takes place in the brain. Indeed, reading this sentence creates long-lasting changes in your brain, because you're going to remember a bit of it tomorrow (I promise you), and this means that the structure of your brain has been modified by this experience. If there was some mental activity that *didn't* change the brain, it would prove Cartesian dualism and would be one of the most amazing discoveries of our time. But this will never happen because Cartesian dualism is mistaken.

By the same token, while the details of mind-brain relationships can be interesting, the fact that the brain is involved in mental life should be seen as obvious—and it sometimes isn't. There was a *New York Times* article, in the Science section, many years ago, titled: "In Pain and Joy of Envy, the Brain May Play a Role." And my reaction was: *May?* Where else is the pain and joy of envy going to be, the big toe?

Some think that neuroscience tells us little about psychology. They are particularly uninterested in findings of localization. The philosopher Jerry

Fodor wrote: "If the mind happens in space at all, it happens somewhere north of the neck. What exactly turns on knowing how far north?"[35] An astute contemporary observer, Matthew Cobb, chimes in with a similar point: "A map—and at their best that is what fMRI data are—does not tell you how something works. Where is not how. The next time you read a claim that a particular ability, or emotion, or concept has been localized to a particular region of the human brain using fMRI, ask yourself, 'So what?'"[36]

I'm not quite so skeptical. While I think that the relevance of neuroscience is often overblown, some of the results really do matter for psychological theory.

Just as one example, in research by Naomi Eisenberger and her colleagues, subjects have their brains scanned while they play a virtual ball-tossing game that they believe is with two other people.[37] Actually, it's a computer program, and it's designed to give them the feeling of being excluded, by having the other characters toss the ball to one another, and leaving the human out.

This hurts. Being shunned is painful, and this study was designed to explore the theory that the pain of rejection shares deep commonalities with actual physical pain. And this is what the brain scanning found: Relative to those subjects who didn't get socially shunned, people in the social exclusion condition had increased activation in parts of the brain such as the dorsal anterior cingulate cortex and anterior insula, the same parts of the brain that are activated when feeling physical pain. This finding has a surprising (though controversial) consequence, which is that interventions that reduce one sort of pain should reduce the other, and indeed there is some evidence that drugs like Tylenol, designed to work on physical aches and pains, can also diminish the hurt of loneliness.[38]

Psychology doesn't reduce to neuroscience, but neuroscience really can tell us interesting things about how the mind works.

We've seen the case for materialism, the theory that the brain is the source of the mental life. We know a lot about how specific activities of the brain correspond to experience and can observe this through tools such as fMRI. Like Santa Claus in the Bing Crosby song, an attentive neuroscientist can tell when you are sleeping and when you are awake—and perhaps isn't that far from telling whether you are naughty or nice, or at least thinking naughty thoughts or nice thoughts.

But there remains what the philosopher David Chalmers has called "the hard problem" of consciousness.[39] How is it that the activities of the brain correspond to conscious experiences? We know that they do; we know that the workings of physical stuff—meat—give rise to what it's like to slam your hand in a car door, or eat scrambled eggs with hot sauce, or kiss someone for the first time. But how does this happen? It seems like magic.

To put the problem in a different way, let's return to computers. My laptop can do smart things like playing chess. This is a rebuke to Descartes, who doubted that mere physical objects are capable of complex action. But as far as we know, my laptop cannot feel the pang of loneliness, the heat of anger, and so on. When it wins or loses, it feels nothing. If I drop it down the stairs, nobody besides me is going to suffer; if I decide to salvage it for parts, this isn't like murder, not even a little bit. So, what's missing? What does one have to add to a machine to give it the capacities to feel? Or can this never happen—must a conscious being be flesh and blood?

There are many avenues to pursue here, and we'll get to some in the next chapter. But I think the honest response, right now at least, is that nobody yet knows.

2

Consciousness

The topic of this chapter can be seen as a quirky diversion. We'll see that behaviorists like B. F. Skinner believed that an adequate science of psychology would say nothing about conscious experience. After all, we don't discuss consciousness when talking about rats, and we're no different from rats. The cognitive psychologists who followed Skinner rejected just about all his views—except for that one. After all, we don't discuss consciousness when talking about computers, and we're no different from computers.

I did my own graduate studies at MIT, which was ground zero for cognitive psychology. My dissertation was on children's language learning, and when struggling with questions like how children figure out what words mean, it never occurred to me to wonder what the experience of learning language felt like for a child. My fellow graduate students and professors studied language, perception, attention, memory, and reasoning, and like me, they thought of these capacities in terms of brain processes and computational mechanisms. Consciousness wasn't relevant. We had philosophers on the other side of campus, in an old World War II timber structure designated as Building 20—let them worry about it.

If we were asked to defend our dismissal of consciousness, we would point out that intelligence does not require sentience. A calculator