Copyright © 2019 By Edge Foundation, Inc. All Rights Reserved.

To arrive at the edge of the world's knowledge, seek out the most complex and sophisticated minds, put them in a room together, and have them ask each other the questions they are asking themselves.

https://www.edge.org/conversation/patricia_s_churchland-the-nature-of-moral-motivation
Printed On Mon November 18th 2019

CONVERSATION : MIND

The Nature of Moral Motivation

A CONVERSATION WITH Patricia S. Churchland [10.16.19]

EdgeCast

Although we have made tremendous progress in understanding many details of the brain, there are huge gaps in our knowledge. What's relevant to me, as somebody who's interested in the nature of moral behavior, is how little we understand about the nature of reasoning, or if I may use a different expression, problem solving. I don't know what reasoning is. For a long time, people seemed to think it was completely separate from emotion, but we know that can't be true.

The nature of problem solving is something that is still very much in the pioneering stages in neuroscience. It's a place where neuroscience and psychology can cooperate to get interesting experimental paradigms so that we can attack the question: How is it that, out of all these constraints and factors, a reasonable decision can be made? That's a tough one. It will require us to find good experimental paradigms and new techniques.

PATRICIA S. CHURCHLAND is professor emeritus of philosophy at the University of California, San Diego, and adjunct professor at the Salk Institute. Her research has centered on the interface between neuroscience and philosophy, with a current focus on the association of morality and the social brain. She is the author of *Conscience: The Origins of Moral Intuition*. **Patricia S. Churchland's *Edge* Bio Page**

THE NATURE OF MORAL MOTIVATION

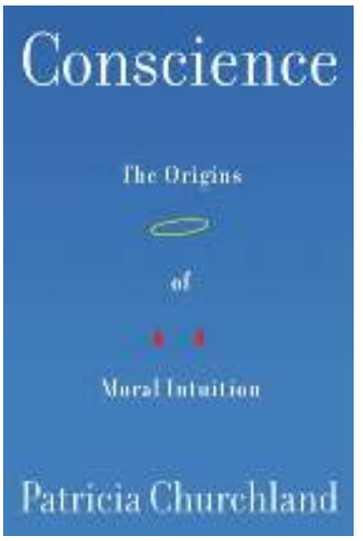

The question that I've been perplexed by for a long time has to do with moral motivation. Where does it come from? Is moral motivation unique to the human animal or are there others? It's clear at this point that moral motivation is part of what we are genetically equipped with, and that we share this with mammals, in general, and birds. In the case of humans, our moral behavior is more complex, which is probably because we have bigger brains. We have more neurons than, say, a chimpanzee, a mouse, or a rat, but we have all the same structures. There is no special structure for morality tucked in there.

Part of what we want to know has to do with the nature of the wiring that supports moral motivation. We know a little bit about it, namely that it involves important neurochemicals like oxytocin and vasopressin. It also involves the hormones that have to do with pleasure, endocannabinoids and the endogenous opioids. That's an important part of the story. The details are by and large missing. And what I would love to know, of course, is much more about the details.

The moral platform provides the basic impetus for sociality, and hence for morality. But that moral platform has to be shaped and directed into very specific kinds of behavior. And for that, we know that the reward system is extremely important. Knowing about norms and practices is also important. What we don't understand well is how exactly they work together so that you know in this condition you should tell a lie, but not in that condition. The decision making of a mammal, especially a human mammal, once they are mature enough to have an understanding of their social and physical world, is under exploration in neuroscience. There's so much we don't yet understand.

What is neurophilosophy? This is a good question, and one I'm asked quite a lot. The very fast answer is it's that space between neuroscience and those big, traditional philosophical questions about the nature of consciousness, about how we learn and remember, about how perception guides our behavior, about how we might think to change the norms within the social group we live in.

Several amazing things happened in the 1970s in neurobiology that made me think much of what philosophers were doing was kind of a waste of time. One of the things that happened came out of Caltech, and that was the split-brain data. The split-brain data was important because, from a philosophical point of view, there had been much argumentation and furrowing of brows over the unity of consciousness. The unity is an essential feature of consciousness, so the argument went.

In California there were a number of patients who had intractable epilepsy, meaning, their epileptic seizures could not be controlled by drugs. Consequently, they had 30 to 40 to 50 seizures a day. Life for these patients was barely livable. The proposal was to split the hemispheres in hopes of diminishing the spread of a seizure from one hemisphere to the other. This was done with several patients. They seemed to recover well. They seemed quite normal. Then, Roger Sperry at Caltech began to test them very carefully. What he noticed was stunning: one hemisphere could know something the other did not, could want something the other did not, could be conscious of something the other was not. This was an extraordinary thing to me at the time because it meant that not only was it *not* a necessary truth that consciousness is unified, but it turns out if you split the brain, you split the mind. I thought that was an amazing thing.

What happened then was that I had the good fortune to be in Winnipeg, Manitoba. I wanted to understand as much as I could of what was known about the nature of the brain, but I ran into a problem. The brain is a three-dimensional structure, and I was looking at two-dimensional pictures. I found the hippocampus on one page, but I couldn't for the life of me see where the heck it was on the other page. So I went down to the medical school one fall afternoon and I went into the anatomy department and told them I needed to understand the anatomy of the brain. The medical school welcomed me. They welcomed me by letting me take any of the courses the medical students are taking. They also gave me a whole human

brain to dissect. The clinicians even asked me to come to grand rounds so that I could see patients who had remarkable and amazing deficits as a result of stroke.

That completely changed how I thought about everything. Suddenly, bickering away about this little distinction and that little distinction in philosophy just didn't cut it anymore. I wanted to know what this thing was, how it worked, how it makes me what I am, how it supports consciousness, and what happens when I sleep. I was consumed by the realization that making progress via science is really fun.

The great fortune of being in Winnipeg, first of all, had to do with the medical school being so welcoming. The second thing had to do with the fact that the philosophy department really didn't care what I did. They didn't put any pressure on me to publish. They didn't care if I ever published anything. They cared that I taught and that I taught well. If I organized myself carefully, it lent me lots of time to do neuroscience. I eventually became attached to a lab that worked on motor control. At the time, I would have rather been in a perception lab, but there wasn't a good perception lab.

So we did motor control, and that changed how I thought about cognition and perception. Rodolfo Llinas, the great neuroscientist, came to visit one day. I complained to him that I wanted to be in a perception lab, and he said, "You're wrong. You're absolutely wrong. We will not understand cognition or the nature of representation in perceptual systems unless we understand how it interfaces with motor control." That changed how I thought about perception and the evolution of the brain.

Rodolfo Llinas' ideas had a huge effect on how I thought about things. Although the time in Winnipeg was amazing, and seeing human patients and what they could and couldn't do was really insightful, we did not have anything like a brain scan. We had X-rays and post-mortem dissections of brains, but when the first PET scans came online, which came out of Sweden, I remember thinking that these were the hoped-for things. This technique would enable us to see where a lesion was so we could make strong correlations between lesions and behavior.

PET scans eventually gave rise to magnetic resonance imaging, and we found a way of using that to look at activity in the brain. Now, of course, it's very normal for people to be put in the scanner and shown a picture so that we can see which parts of their brain light up. The change over those fifty years that was made possible by scanning was absolutely tremendous. But that was only part of it. Many other techniques came online as well, techniques at every level, from cells to systems to the whole brain. Neuroscience had suddenly become this tremendous wave that we were all trying to catch. There are still techniques that we desperately need, and that will come, to enable us to address the nature of what it is to be conscious, and determine how the brain supports consciousness by using these techniques, but we will have to wait.

Mike Gazzaniga, who was part of Sperry's lab at the time, saw how absolutely important these split-brain studies were. He was then at Dartmouth, where there was another set of patients and much further and more developed work was done with split-brain patients. Mike had a huge impact on me because he was not just interested in collecting fascinating data; he saw the big questions. He knew this was going to have a lot to do with the nature of how we understood the mind itself.

The other astonishing discovery in the 70s, which we are still working through, had to do with the surgical subject, H.M. H.M. had intractable epilepsy, and the suggestion was that the hippocampal structures on both sides be removed, which was done. His epilepsy seemed cured, but then he could not learn anything new. He had a thirty-second range of memory. If he was interrupted or distracted, you'd have to explain it to him all over again. Every day when the doctor came in, he had to be reintroduced to him. This told us that there was something very special about hippocampal structures and memory for events, episodes, and facts. It was unexpected that it was separate from other structures that had to do with learning skills and routines. H.M. was able to learn a skill, but he couldn't remember having learned it.

This, like the split-brain data, set in motion a huge research enterprise into the nature of learning and memory. For one of the first times ever, we were asking how many learning

systems there are. How many different mechanisms are there for learning things? How is recall ever possible? Notice we still have not answered that question. When I ask you if you remember your first kiss, I'm betting that within a hundred milliseconds you do. How is that possible? We don't know.

When David Chalmers talks about the "hard problem," I think, well, there are many hard problems about the brain. The research that came out of this initial discovery, tragic though it was, that the hippocampus is essential in some way for learning new things, changed neuroscience profoundly. Sometimes when people ask me if I think there might be a soul that survives our bodily death, I say, "Look, you need the hippocampus to learn anything new. When your brain dies, your hippocampus goes too. It is done. So what's your soul going to use for information? How is it going to learn anything?"

This work on H.M. was also related to something else that was profound. There was a patient of the Damasio's, who was like H.M. except worse. This patient's name was Boswell. Boswell not only couldn't learn anything new, but he had lost essentially all autobiographical memory. Boswell knew he was born in Iowa, but that was it. So when I asked him if he was ever married, he said, "It's hard to say." But notice how he responds, right? He's gracious, trying, and he knows in a certain sense that he doesn't know. This is a man with no autobiographical memory, but still has the capacity to use his social skills. You wouldn't know for the first three minutes of having met him that there was anything wrong.

For me as a philosopher, this was an incredible result, and I'll tell you why. In the 18th century, John Locke had the idea, which was not a bad idea, that our sense of self, our very selfness itself, depends on autobiographical memory. Without autobiographical memory, we wouldn't have a self. Most philosophers thought that a) Locke was right, and b) autobiographical memory was a necessary condition for having a sense of self. But my argument was, Boswell clearly has a sense of himself. When he's introduced to me, he shakes my hand, he asks me where I'm from and how I like the weather. He has a sense of self, albeit a diminished one. He doesn't know that he was married, can't recognize himself or his children in a picture, but he doesn't lack all sense of self. Basically, the philosophers ignored me. They wanted nothing to do with it.

I did start as a philosopher. When I was an undergraduate, we studied Aristotle, Hume, and Descartes. These guys were trying to understand phenomena; they weren't just playing around. But by the time I got to graduate school, it was quite clear that people who were doing what they called philosophy of mind were no longer interested in phenomena; they were interested in words. What do certain words really mean? What does the word "knowledge" really mean? Hume wasn't interested in that so much as, how is it that we know anything?

By the time I was well into graduate school, I had begun to sour because I thought, it's all well to talk about what words really mean, but ultimately what you want to know is about the phenomenon itself. I was a bad graduate student because I didn't want to play that game. By the time I was an assistant professor in Winnipeg, Manitoba, I was ready to kick over the traces.

How is philosophy different from science? The first time I felt like I had an insight to that question came when I was a graduate student at Pittsburgh and read Willard van Orman Quine. Quine said philosophy and science are continuous, that philosophers may be asking questions where we can't yet get data, but when we can get data, those questions will be considered scientific questions. There may be some parts of philosophy, such as aesthetics, where it's not related to science. Philosophy of mind, how we work, what it is to have consciousness, what it is to think, how perception and thinking and motor control integrate —those are questions you can't answer by sitting in an armchair and dreaming up thought experiments.

What happened in the early 80s was one of these grand chance events. I was asked to give a talk at Johns Hopkins. I had read in *Nature* that there was this computational neuroscience guy at Johns Hopkins called Terry Sejnowski, and I thought, I've got to meet this guy and figure out what he knows. So off I went, and I met Terry.

The other thing that happened was that Francis Crick was also there, and he also gave a

talk. After I gave my talk, which was about needing to understand the brain in order to understand the mind, he was clearly taken aback. He said, "I've been waiting for years for philosophers to say this!" He was so excited by the idea that a philosopher was going to answer questions about the nature of ourselves. We had hours of conversation, Francis and I. After that, he organized things so that Paul and I ended up moving to San Diego. It was a wonderful thing. San Diego had been the Mecca of neuroscience for me. There was Ted Bullock, who had done comparative neuroanatomy and neurophysiology, just one of the great archbishops of science. There was Al Selverston, who had done the first model of a piece of behavior where we knew the wiring. There was Larry Squire, who was working on declarative memory in H.M. I was in heaven.

One of the first things Francis did was arrange for me to be an adjunct professor at the Salk Institute, which was great. We had lunch one day, and he told me we needed to get Sejnowski here because he was the kind of person who is going to lead in getting real theoretical insight into the nature of the nervous system. So we went to work on it, and low and behold, Terry got an offer from Salk that he couldn't refuse.

Terry's lab had tea every afternoon at about 3:30. Francis always came for tea, his postdocs came for tea, and it was the most lively, outrageous, no-holds-barred philosophy seminar I've ever been in. Some graduate student would ask, "What the heck is an explanation anyway?" and off we'd go. Francis of course participated in those things because he loved that kind of rough and tumble. That first decade or so of being at the Salk was wonderful for me.

Francis didn't just know about the history of microbiology, he knew how to put it in the context of a much longer view of the history of science. It was wonderful to talk to him about the history of molecular biology and what it was like before they knew there were two kinds of RNA—messenger RNA and transfer RNA. There was a time when they just knew there was RNA, and these guys got these weird results, and Francis said, "So we decided to do the intelligent thing and solve this." And they did. I think it was Sydney Brenner who said, "Wait, maybe there are two kinds of RNA. Let's call it messenger and transfer." Suddenly, that hypothesis meant you could explain all of the disparities. What a cool thing.

I came to know Jim Watson partly through Francis because Jim would come to California from time to time. On one occasion Jim asked me to give a talk at Cold Spring Harbor, which I did. I didn't know him as well as I knew Francis, but he was interested in many detailed questions of molecular biology. Francis had shifted, as he was fond of saying, "I don't really go to molecular biology talks anymore because I can't follow them." What he meant was that there had been so much progress, and he couldn't follow it because he was spending his time learning about brain. And that was what he and Terry and I had in common.

I knew when I was in Manitoba that I had a book in me about the relation between neuroscience and philosophy. I knew it couldn't be just a short book, and that I had to be able to explain at least enough of the history of philosophy and enough of the history of neuroscience to see why now is a propitious time for these to come together. It wasn't quite finished when we moved to San Diego, and I will say Francis was a tremendous help in pushing me along.

Eventually in 1986, the book came out and he was tremendously supportive. The book got a very bad review from (neuroscientist) Gunther Stent, who worked on the leech. It was in *Nature*. I was heartbroken. I was not mature enough to realize it didn't matter. I asked Francis if he thought I should respond, and he said, "Absolutely not. Say nothing." He was right, of course, that it didn't matter. It was very helpful to get advice from him. He was a very practical person. He often doesn't seem like that because he would have these very rich, insightful far-seeing ideas, but he was a very practical person.

The impact of my idea that there should be neurophilosophy is a little hard to assess. When the book came out, many students who were in philosophy and had thought philosophy would help them understand the nature of the mind changed and went into neuroscience. There has never been a talk I have given anywhere in the world where I haven't had faculty and postdocs and others tell me they quit philosophy because they read neurophilosophy.

Part of what happened then was, and I think it probably would have happened anyway, people could see that neuroscience was where the action was. It wasn't in analyzing words and making things clear. Graduate students and postdocs from physics, from biology, from math went in droves.

Part of the impact, as a result of the neurophilosophical approach, was that they could see, if not in detail at least in outline, the links between brain and mind. Thinking is just a function of the brain, and someday we'll understand what that really is. What I didn't have (among philosophers), apart from maybe one or two people within philosophy, was people saying, "Oh my God, neurophilosophy is the way to go." One person tried to start a journal called *Neurophilosophy*, out of MIT, and MIT Press was very keen on the idea, but some well-placed, distinguished philosopher commented that it would be too associated with Pat Churchland. They couldn't have that. The philosophers were very mean spirited about many of these things. I had one distinguished philosopher tell me I had no right to the word neurophilosophy. Whose word was it anyways? It's nobody's word.

John Bickle was one graduate student from North Carolina who picked up on it and now does work with Alsino Silver, who's a neuroscientist. There were others as well. We can see the old guard is still doing good old-fashioned word-meaning analysis. The younger people are interested in the overlap. They want to know how neuroscience is affecting how we understand the nature of thought, the nature of having a self, the nature of making decisions. Mike Gazzaniga said to me one day, "There are so many philosophers doing history of philosophy. They're reading Hegel, and Fichte, and Nietzsche, and Plato. But do you think that if Aristotle were alive today, he'd be reading Aristotle?"

This was a funny way to ask a very interesting question. I had to say no. Aristotle would be doing today what he was doing then, trying to find out how things work. He would be asking questions like, "How many teeth do we have? Are baboons similar to chimpanzees? What does the heart really do?" These were things he was interested in. Of course, he didn't always get the right answers. Needless to say, if he were alive right now, he'd be working on the brain.

Although we have made tremendous progress in understanding many details of the brain, there are huge gaps in our knowledge. What's relevant to me, as somebody who's interested in the nature of moral behavior, is how little we understand about the nature of reasoning, or if I may use a different expression, problem solving. I don't know what reasoning is. For a long time, people seemed to think it was completely separate from emotion, but we know that can't be true.

The nature of problem solving is something that is still very much in the pioneering stages in neuroscience. It's a place where neuroscience and psychology can cooperate to get interesting experimental paradigms so that we can attack the question: How is it that, out of all these constraints and factors, a reasonable decision can be made? That's a tough one. It will require us to find good experimental paradigms and new techniques. There are psychologists who address these questions about problem solving. And in my own deep-seated skeptical way, I'm always out there kicking the tires.

There are many things about Danny Kahneman's approach that I find very appealing, this idea that there are two systems—one that uses logic and one that doesn't—but from the point of view of neuroscience, it doesn't really add up. This is partly because there are many actions that we do very quickly if we have the right skills that are extremely intelligent. Think of a quarterback in football, think of a hockey player, a neurosurgeon, an eye surgeon. Automatic? Well, no, not exactly. If something weird happens, he's responsive.

This dichotomy is okay as a first pass, but it doesn't pass muster. It's not going to work. Gerd Gigerenzer has ideas about heuristics that I find appealing in many ways. We need to find ways of testing neurobiologically. Can we find in the nervous system of a rodent, evidence that a heuristic, a rule of thumb that other things being equal, the rodent uses. It's quite possible that we'll find some evidence for it, but that remains to be seen.

As a neurophilosopher, the question of consciousness is quite interesting. The way I've always thought about this is quite practical. Let's start where we can get some traction. We want to know the difference between a brain that's awake and a brain that's in deep sleep.

We want to know the difference between a brain that is in coma and a brain that is awake and functioning, or a brain that is anesthetized and one that isn't. We want to know the difference between paying attention to something and not.

This is a way to get into the problem and begin to find out what the circuitry is that supports awareness, because there's no doubt that there is a difference between being in a coma and being awake. There are some data relevant to all of those questions. There are data relevant to sleep and dreaming, for example. Stan Dehaene has lovely stuff. Giulio Tononi has lovely stuff. It's a small baby step forward. Nick Franks has important data on what happens under anesthesia. It looks like the inhibitory cells are upregulated and the excitatory cells are downregulated.

But what you want to know is why that leads to conscious disappearance. Why do we lose consciousness under anesthesia? There are data concerning the nature of attention, both top-down attention and bottom-up attention. When a dog suddenly barks and I turn, that's one kind of attention. When I begin to think about how to get from your office to Columbia, that's a different kind of attention. So we have bits and pieces of all of that.

What some people say is, yeah, but you'll never get the answer. How do they know that? It's like my old biology teacher on the farm saying, "We'll never understand the nature of life because living molecules are different from dead molecules." Really? Never is a long time, especially when you've got smart postdocs and graduate students. You can say we don't understand it now, but why say we will never understand? For a certain subset of philosophers, it's very self-serving. It allows them to say, "I can tell you about the necessary features of consciousness because that's what philosophers do." The answers will come, just as they did with DNA. The answers will come. Saying that they aren't here now so they'll never come—this is not rational.

What's Related

People

Patricia S. Churchland

Philosopher and Neuroscientist; Author, Conscience: The...

Video

The Nature of Moral Motivation

Beyond Edge

Patricia Churchland's Website

Books

Conscience: The Origins of Moral Intuition
By **Patricia S. Churchland** Hardcover [2019]

Touching a Nerve: The Self as Brain
By **Patricia S. Churchland** Hardcover [2013]

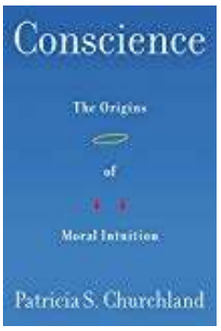

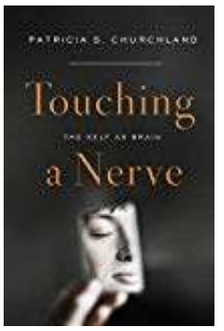

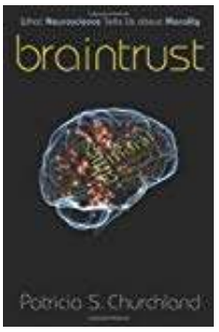

Braintrust: What Neuroscience Tells Us About...
By **Patricia S. Churchland** Hardcover [2011]

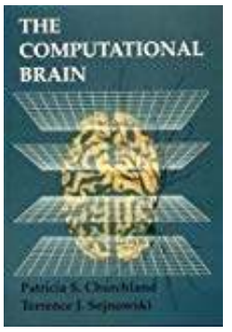

The Computational Brain (Computational...
By **Patricia S. Churchland**
Paperback [1994]

Tags

conscience
consciousness
moral motivation
neurophilosophy
neuroscience
philosophy

John Brockman, Editor and Publisher
Russell Weinberger, Associate Publisher
Nina Stegeman, Associate Editor

Contact Info:editor@edge.org

Edge.org is a nonprofit private operating foundation under Section 501(c)(3) of the Internal Revenue Code.
Copyright © 2019 By Edge Foundation, Inc All Rights Reserved.