

Class #18 - The Mind

I. Distribute Papers to Peer-Review Groups

II. From Self to Mind

In our unit on personal identity, we were looking for either a reductive theory of the self or an argument for why such a reduction is impossible.

A reduction is an explanation in simpler or more fundamental terms.

Now we are starting a unit on the philosophy of mind.

The central problem in the philosophy of mind is to determine the what minds and mental states are.

Is the mind a non-physical substance that thinks?

How are minds related to bodies?

Are minds explicable in physical terms, as behaviors or brains?

Are minds the software of the brain?

What is consciousness, and how do conscious states relate to other mental states?

Again, we want a reduction of mental states or an explanation of why no such reduction is possible.

Questions about the nature of mind are thus similar, in form, to questions about the self.

There are similarities in content as well.

Descartes's theory of the self is essentially the same as his theory of mind: the self is the mind is the soul.

Note that Descartes has a substantial view of the self: the self is identified with a substance, the soul.

The body theory, even, perhaps, the biological theory, is also a substantial theory of the self.

But at least one promising theory of the self, Locke's view that we identify the self with consciousness, is not a substantial view.

Locke argues that the self is not a thing, or substance, but a connection among mental states.

We called Locke's view both the memory theory and the consciousness theory.

Locke seems to think of memory as a kind of consciousness.

Memory is like storage for conscious experiences, which we can bring to mind when we recollect.

Reid pointed out that a memory of an experience is in many ways not like the conscious experience itself.

Mr. Locke attributes to consciousness the conviction we have of our past actions, as if a man may now be conscious of what he did twenty years ago. It is impossible to understand the meaning of this, unless by consciousness be meant memory, the only faculty by which we have an immediate knowledge of our past actions. Sometimes, in popular discourse, a man says he is conscious that he did such a thing, meaning that he distinctly remembers that he did it... But this ought to be avoided in philosophy, otherwise we confound the different powers of the mind, and ascribe to one what really belongs to another. If a man can be conscious of what he did twenty years or twenty minutes ago, there is no use for memory, nor ought we allow that there is any such faculty. The faculties of consciousness and memory are chiefly distinguished by this, that the first is an immediate knowledge of the present, the second an immediate knowledge of the past (Reid 347).

Reid's criticism of Locke raises important questions about mental states.

If there is a self, and if we are not going to take it as a brute fact, then we need to have a more subtle understanding of the nature of mental states, of conscious experience.

III. Two Kinds of Mental States

There are two general approaches to characterizing mental states.

These approaches correspond to two different kinds of mental states.

The first approach focuses on consciousness.

My conscious experiences include my sensations and my emotions.

I can be conscious of the dog barking in the yard, or the smell of garlic, or the blue sky, or a tickling sensation.

I can be conscious of my joy, or my anger, or my anxiety, or my appreciation.

I can be conscious of my belief that it is raining or that the tangent to a circle meets a radius of that circle at right angles.

Descartes believes that all mental states are conscious and that the only conscious things are minds.

One problem with Descartes's view is that there seem to be unconscious mental states.

I might have unconscious emotions, feelings, or attitudes.

My non-occurrent beliefs seem to be unconscious, as well.

I believe that the square root of two is irrational, even when I'm not thinking about it.

Interest in mental states like beliefs leads to the second approach to characterizing mental states, via intentions.

Intentions are mental states that have some content, that are about something, that represent.

They are sometimes called propositional attitudes, since they are naturally taken to be relations between people and a proposition.

A proposition is a mind- and language-independent fact: that snow is white, that $2+2=4$, that I will eat lasagne for dinner.

Belief is the paradigm intentional state.

I can believe that snow is white or desire that I will eat lasagne for dinner.

As well as having intentions, we attribute them to other people.

It is an open question whether all mental states are of one kind or another, or whether conscious experience and intentionality are two different kinds of mental states.

IV. Consciousness, Minds, and Persons

The first theory of the mind we will study come from Descartes.

Then we will look at more-recent work on the philosophy of mind.

Until the twentieth century, few philosophers took the possibility of a physical theory of mind seriously.

Over the last century, philosophers of mind developed a variety of theories attempting to accommodate a materialist framework.

We will look at four distinct theories of mind: dualism, behaviorism, identity theory, functionalism.

At the end of our unit, we will look specifically at the problem of consciousness.

There are (at least) two problems of consciousness, an easy problem and a hard problem.

The so-called easy problem involves determining the functions of the brain.

Neuroscience is essential for solving the easy problem.

For example, if we want to know about how we focus our attention on, say, the tree outside the window, we have to know about perceptual systems, attention spans, and phenomena like staring.

An understanding of the brain and the rest of the body will be part of any decent scientific or philosophical explanation of attention and focus.

The hard problem is to explain the connection between brains and conscious awareness. It arises from recognizing that figuring out the neural correlates of consciousness does not suffice for explaining what it is to be conscious. The fact that consciousness involves experience, rather than function, makes the problem difficult.

We don't know whether cognitive neuroscience can tell us anything, or everything, about who we are. On the one hand, it seems obvious that a complete description of our bodies, [especially our brains](#), will suffice to explain our minds, and thus who we are. On the other hand, the nature of conscious awareness seems to resist physical explanation. If we had physical explanations of consciousness, then, at least in theory, we could construct machines that think (and not just by procreating!) But the idea that a physical machine could think is uncomfortable, for many of us. Consider a standard objection to a materialist theory of mind, cited by Alan Turing as a criticism of his defense of artificial intelligence.

Not until a machine can write a sonnet or compose a concerto because of thoughts and emotions felt, and not by the chance fall of symbols, could we agree that machine equals brain, that is, not only write it but know that it had written it. No mechanism could feel (and not merely signal, an easy contrivance) pleasure at its successes, grief when its valves fuse, be warmed by flattery, be made miserable by its mistakes, be charmed by sex, be angry or depressed when it cannot get what it wants (Alan Turing, "Computing Machinery and Intelligence," *Mind*, 1950).

It is easy enough to make machines produce [poetry](#) or [art](#) or [music](#). Computers can defeat us in [chess matches](#) and on [Jeopardy](#). Machines are better than us at lots of menial tasks. But, for a long time, we were able to distinguish ourselves from mere machines with the observation that the abilities of machines to perform even complicated tasks were due to our intelligence, our minds, and not their own. Machines can only do what we tell them to do. The objection above is not that we can distinguish ourselves by what we do or make. It is that real intelligence involves internal processes that cause those behaviors or products.

Of course, the internal processes of machines may be unobservable. We can see levers and dials and circuits. But we don't know what it would mean to see the intelligence of a machine. How does one see a mind? We don't even know what it would mean to see the intelligence of another person. We infer the existence of mental states in each other from our interactions with each other's behavior, especially our language.

Many science fiction stories (including the excellent movie *Blade Runner*) involve sentient robots, machines built by humans which seem to think and feel and be conscious. We may, in the future, have to decide whether to see some such machines as persons. One important question about whether to see them as persons is whether they are producing their behaviors as a result of conscious, inner experience, or as a result of the "chance fall of symbols." We can easily make artifacts which look like humans, but how they look is immaterial. The question of memories will play a role. Some robots could have memories implanted; then again, so can we.

If some such machines are persons, then we can conclude mental states will seem to be reducible, in some ways, to physical states.

Some contemporary materialists hold this view.

Some even talk as if experience is completely explicable in neural terms.

Eliminative materialists like Patricia and Paul Churchland believe that our ordinary language will be abandoned in the future for a more precise language about our brains and bodies.

For instance, both [Paul] and Pat [Churchland] like to speculate about a day when whole chunks of English, especially the bits that constitute folk psychology, are replaced by scientific words that call a thing by its proper name rather than some outworn metaphor. Surely this will happen, they think, and as people learn to speak differently they will learn to experience differently, and sooner or later even their most private introspections will be affected. Already Paul feels pain differently than he used to: when he cuts himself shaving now he feels not “pain” but something more complicated—first the sharp, superficial A-delta-fibre pain, and then, a couple of seconds later, the sickening, deeper feeling of C-fibre pain that lingers. The new words, far from being reductive or dry, have enhanced his sensations, he feels, as an oenophile’s complex vocabulary enhances the taste of wine. Paul and Pat, realizing that the revolutionary neuroscience they dream of is still in its infancy, are nonetheless already preparing themselves for this future, making the appropriate adjustments in their everyday conversation. One afternoon recently, Paul says, he was home making dinner when Pat burst in the door, having come straight from a frustrating faculty meeting. “She said, ‘Paul, don’t speak to me, my serotonin levels have hit bottom, my brain is awash in glucocorticoids, my blood vessels are full of adrenaline, and if it weren’t for my endogenous opiates I’d have driven the car into a tree on the way home. My dopamine levels need lifting. Pour me a Chardonnay, and I’ll be down in a minute.’” Paul and Pat have noticed that it is not just they who talk this way - their students now talk of psychopharmacology as comfortably as of food ([The New Yorker profile of the Churchlands](#), February 12, 2007).

Bodies, the Churchlands argue, are nothing more than information-processing machines.

The question of whether some machines are persons will guide the rest of the course.

First, we will look at the nature of mental states and conscious experience.

Then, we will look at the moral question of what makes a person.

A large part of the answer to that question involves the nature of consciousness.

There may be other factors as well.

V. Bodies, Minds, and Animals

Regarding the nature of mind and its relation to the possibility of artificial intelligence, we can distinguish between a liberal and a conservative view.

On the liberal view, minds are just information processors.

Defenders of artificial intelligence hold the liberal view and claim that machines have minds.

Consider IBM’s chess-playing computer which beat Garry Kasparov in 1997.

Saying Deep Blue doesn’t really think about chess is like saying an airplane doesn’t really fly because it doesn’t flap its wings ([Drew McDermott](#)).

In contrast, on a chauvinistic view only things built like us humans have minds.

Robots, for example, are just machines with no appreciable or important inner states.

On an even narrower view, called solipsism, I have good reasons only to believe that I have a mind.

A related, interesting question, one that we might keep in mind as we proceed, is whether animals have mental states.

There is behavioral evidence of some higher-level functionality in some animals, including some surprising if rudimentary language use.

[Smart Chimps](#); [Painting Elephant](#)

What we say about the nature of mental states will be general.

It will apply to all sorts of things: humans, robots, aliens, and animals.