

Who Is Sitting Behind Your Eyes?

The self, the empty boat, and the science of "I"

Learning Psychiatry YouTube episode — Reading Edition

By Said Fard

There is an old story about a man crossing a river.

It is early. The water is flat and grey, and the mist has not lifted yet. He is rowing a small boat, and he is not in a good mood. Something happened before he left — an argument, a bad night, it does not matter. His shoulders ache. His hands are cold on the oars. He wants to get to the other side and be left alone.

And then something hits him.

Another boat comes out of the mist and strikes the side of his boat, hard enough to throw him sideways. Water comes over the edge. His knee slams into the wood. And before he has even finished falling, something rises in him — fast, hot, and completely certain. He knows exactly what has happened. Someone was not looking. Someone was careless. Someone did this to him.

He pulls himself up, turns around, and opens his mouth to shout at whoever is steering that boat.

But there is nobody in it. The boat is empty. It came loose somewhere upstream, and the river carried it down, and it struck him the way a branch or a wave would have struck him.

Now watch what happens inside him, because this is the strange part. Nothing about the collision has changed. His knee still hurts. The water is still in the bottom of his boat. He is still wet, still cold, still late. The damage is exactly the same as it was five seconds ago. But the anger has nowhere to go. It was not really about the water or the knee. It was about the person he was sure was in that boat. And when the person turns out not to exist, the anger loses its shape and drains away, and he is just a man on a river with a wet knee.

So here is the question that story leaves behind, and it is not the question it looks like.

When we are angry at someone, are we angry at what happened — or are we angry at the person we imagine sitting behind their eyes, deciding to do it? Because we do imagine that. We imagine that behind every face there is someone in there. Someone steering.

And if we imagine a hidden chooser inside other people, we imagine one inside ourselves even more strongly.

I'm Said Fard, a psychiatrist in Sweden. This is a question I have never quite been able to put down — not in the clinic, and not at four in the morning when I cannot sleep.

Who is sitting behind your eyes?

Chapter 1 — The Feeling of Being Me

Before we take the self apart, we should be honest about how solid it feels. The self is not just a theory philosophers argue about. It is something you are experiencing right now, without trying.

Think about the morning. You wake up, and you do not have to work anything out. You do not check who you are, or prove it, or look for evidence. It is simply given. The world is already there, seen from one particular place — your place. The body in the bed is yours. The thoughts that start moving are yours. And it feels, with no effort at all, as if there is someone in here receiving all of it. Someone behind the eyes, looking out.

That feeling is one of the most convincing experiences we have. It is also one of the hardest things to find when we go looking for it. Philosophers have looked. So have people who meditate, psychologists, and people with brain scanners. None has found a small person in there. What they find is something stranger, and much more interesting.

Start with the simplest part. Your experience has a direction. It comes to you from one particular place — the place where you are. If you are walking, the street opens in front of you and closes behind you. If you are washing dishes, the sink is here and the window is over there, and both are arranged around a centre: you. Philosophers call this the first-person perspective. In ordinary language, it means that you do not simply see the world. You see it from somewhere. A camera on the wall could record the same room, but the room would not be organised around the camera in the way it is organised around you.

That is the first layer. Here is the second.

Lift your hand. Not far — just enough to notice it. Now, what did you actually experience? A hand moved, but that is not how it felt. The movement came with two feelings at once, and we need to separate them for the rest of this episode.

The first is that the hand is yours. Not a hand — your hand. You did not have to check, compare it with a photograph, or look for a scar you recognise. The sense that the hand belongs to you came for free, built into the experience. Researchers call this the sense of ownership: the feeling that this body, this movement, this thought is mine.

The second is that you did it. The hand did not rise on its own; you raised it. That is a different feeling: the sense of agency, the feeling of being the one who causes an action.

Most of the time, ownership and agency arrive together, and we do not notice that they are different. But they can come apart. If someone takes your arm and lifts it for you, the arm is still unmistakably yours — ownership is intact — but you are not the one moving it. Agency is gone. Think also of a muscle twitch, a hiccup, or a startled jump when a door slams. Your body. Not your doing.

So already, in something as small as lifting a hand, there are separate parts. And there are more.

There is also the sense that things are happening to someone. Not just pain, but pain happening to you. Not just a memory of a summer, but a summer that happened to you and not to somebody else. Experience seems to arrive as experience for someone.

And then there is time. You feel like the same person who woke up this morning, and the same person who was seventeen. The photographs may look like a stranger. Your opinions and your body may have changed enormously. And still, something says: that was me. Same one. Continuous.

Philosophers have names for these layers, and one distinction matters here.

One influential description comes from the philosopher Shaun Gallagher. In 2000, he distinguished two senses of the word “self”. The minimal self is the immediate sense of being the subject of an experience, here and now, with no life story attached. You do not need to remember your childhood to feel that this is happening to you. The narrative self is the self built across memory and time, the one with a history. We will return to that later. The philosopher Dan Zahavi develops a related idea: experience is not simply there. It has a quiet sense of being for me, built into the experience rather than added afterwards by thought.

These ideas are useful, but the minimal self is not something anyone found under a microscope. It is a framework for describing experience, and philosophers disagree about it. Some think this basic sense of experience being for me is real and fundamental. Others think several different experiences have been bundled together under one label. The philosopher Marie Guillot has argued this in detail. So when I say “minimal self”, hear it as a useful proposal, not as a discovery.

What is not in dispute is the experience itself. There is a point of view. The body feels like mine. The actions feel like mine. The thoughts feel like they are happening to someone. And that someone feels like it has been here a long time.

Which brings us to the awkward part.

Each item in that list carries a sense of “mine”. This body is mine. This movement is mine. This life is mine. But if all of it is mine — where, exactly, is the owner? The man on the river turned around expecting to find somebody steering. When we say “my thoughts”, we make a similar assumption. We imagine somebody in the boat.

So let us turn around and look.

Chapter 2 — The Self We Cannot Find

The strange thing about the self is that it feels obvious while we are living it, but becomes hard to locate when we look for it. Thinkers in very different traditions found different parts of this problem, and none solved it cleanly.

Start with the most famous.

In the seventeenth century, René Descartes tried to discover what could not be doubted. He questioned his senses, the external world, even whether a deceiver was feeding him a false picture. But he could not doubt that doubting was happening.

In the 1637 Discourse on the Method, he wrote the famous “I think, therefore I am.” Four years later, in the Meditations, the wording is more careful: “I am, I exist” —

necessarily true whenever it is thought. The argument secures existence in the moment of thinking. Descartes believed in much more, but the argument itself does not prove a permanent inner person with a lifetime attached.

But Descartes was not the first to imagine a self that could remain even when the world and the body were stripped away.

More than five centuries earlier, the Persian philosopher Ibn Sina — known in the West as Avicenna — proposed one of the most elegant thought experiments in the history of philosophy. It is often called the Floating Man.

Imagine, he says, a human being created all at once, fully formed, but suspended in empty air.

He sees nothing.

He hears nothing.

He touches nothing.

His limbs are separated so that no part of his body touches another.

Now ask: would he still know that he exists?

Avicenna's answer is yes.

Even without seeing his body, touching his body, or receiving ordinary sensory information from the world, the floating man would still be aware of himself. He would not know his height, his face, his history, or his place in the world. But he would know that he is.

This is philosophy, not neuroscience or experiment. But it matters here because it pulls in the opposite direction from many modern ideas in which the body is central to the sense of self. For Avicenna, self-awareness is not built up from the body. It is more basic than bodily perception. The self is immediately present to itself.

That creates a clear tension.

Descartes says: even if everything is doubted, thinking reveals existence.

Avicenna says: even without the sensed body, self-awareness remains.

Both point toward a self that seems immediately given — something that does not need to be found in the world because it is already present before the world is examined.

Then came David Hume, who did something different.

He went looking for it.

Hume took the opposite route. If there is a self, he thought, I should be able to find it by looking inward. What he found were particular perceptions — heat and cold, pleasure and pain, thoughts and sensations — but never a bare owner standing behind them.

So he suggested that perhaps the self is not a separate thing that has experiences, but a bundle of experiences in constant motion.

Then Hume ran into a problem. In an appendix to the Treatise, he admitted he could not explain how separate perceptions become one connected life. If there is no owner and no perceived connection binding them, why does experience arrive as a coherent stream rather than scattered fragments?

He had no clean answer. Removing the owner did not automatically explain what the owner seemed to explain.

There is one tempting escape: perhaps some part of the brain is the real observer, watching the images, reading the thoughts, deciding what to do.

This creates the homunculus problem — the little person inside. If the inner observer has an experience, we still have to explain who or what observes that observer. Put another observer inside and the same question returns. We have not solved the problem; we have only moved it one floor down.

So we are left in an uncomfortable position, and it is worth leaving it unresolved for a moment.

We cannot easily live without the word “I”. Try removing it from one day — from your plans, your regrets, your relationships, or the sentence you would use to explain why you are tired. And yet when we look inward for the thing that word points to, we come back with what Hume found: perceptions, sensations, thoughts, a great deal of activity — and no one standing behind it.

But if Avicenna is right, there is something about self-awareness that does not wait for the body or the world. If Hume is right, when we inspect experience, we never find that self as a thing. And if the homunculus argument is right, putting a little observer inside the mind only repeats the problem.

That raises another possibility: perhaps we have been asking the wrong kind of question. Perhaps we are looking for a thing when what is there may not be a thing.

Long before Descartes and Hume, and in a different way from Avicenna, other traditions had reached something close to that conclusion — and built entire practices around it.

Chapter 3 — The Empty Boat and the Chariot

Let us go back to the river.

The empty-boat story is Chinese and more than two thousand years old. It comes from the Zhuangzi, a foundational Daoist text, in the chapter usually called “The Mountain Tree”. It is not a Buddhist sutra, although Zen and later Buddhist teachers adopted it.

I changed the story too. In the Zhuangzi, an empty boat and an occupied boat are contrasted; the man does not first become angry and then discover that the boat is empty. At the start of this episode, I turned that comparison into a sequence and asked a different question: what changes when we discover there is nobody in the other boat? That reading is mine, not Zhuangzi's.

The collision and the pain are unchanged. What disappears is the agent the man had assumed was behind them. I am not claiming that belief in an inner owner causes anger. The narrower point is that anger often assumes a chooser — and we make that

assumption about other people and about ourselves.

Elsewhere in Asia, another tradition examined that same assumption more systematically.

The key word is *anattā*. It is usually translated as not-self, and it is one of the most misunderstood ideas in the world.

Not-self does not mean that people are unreal. Your mother is still there. Your actions still matter. The Buddhist texts explicitly reject the opposite view as an error. The question is narrower: behind all the changing activity of a life, is there a permanent, independent thing that owns it and controls it?

The texts approach this more like a method than a fixed statement. You take the whole of a person and break it down into what is actually there.

Traditionally, this comes under five headings. There is the body — physical form. There is feeling-tone: the raw pleasantness or unpleasantness of an experience, before we think about it. There is perception, which recognises things: that is a face, that is a threat, that is my name. There are mental formations, a broad category that includes intentions, habits, impulses, and the push towards action. And there is consciousness, the basic awareness that any of this is happening.

That is abstract, so let us put it on the river.

The man's boat is struck. What is actually there, in that second?

There is the body: a jolt through the knee, a lurch, heat rising in the chest, hands gripping. There is feeling-tone: this is unpleasant, immediately and without argument. There is perception: someone hit me; someone was careless. There is the impulse: turn around, shout. And there is awareness of the whole thing happening.

Look at that list and ask the question the texts ask. Where is the one who is angry?

You will not find him among the parts. You will find a body reacting, a tone of unpleasantness, an interpretation, an urge, and awareness. Each arrives, does something, and passes. What you will not find is a sixth thing standing behind the other five, receiving them and deciding. And yet, in the moment, it feels overwhelmingly as if there is one — which is exactly what makes the man turn around expecting a face.

This is where the chariot image becomes useful. In an early Buddhist text, the nun Vajirā compares a person to a chariot. Put the parts together in the right way and the word “chariot” works. You can ride it, sell it, or crash it. But if you search for a chariot as an additional thing beyond the organised parts, you will not find one.

And that does not mean chariots are unreal. That is the hinge: something can be real without being an extra object in addition to its parts. It can be real as an arrangement, a pattern, a functioning whole.

Scholars still disagree about how far Buddhist not-self goes — whether it is a claim about what ultimately exists or mainly a practical strategy for loosening attachment. We do not need to settle that debate here. For our question, the important move is the third option: the self need not be either a hidden object or a fake.

For our purposes, the practical shape is clear enough, and it lines up with where Hume left us.

Hume went looking for the owner and found only perceptions. He later admitted that his own account could not explain what held those perceptions together. The chariot offers a different answer: of course you found only perceptions. Those are the parts. The mistake was expecting an extra thing behind them.

If that is right, then philosophy has been circling this problem for more than two thousand years, and modern science eventually approached it from the other direction. In recent decades, researchers stopped arguing about the owner and started measuring the parts.

Chapter 4 — The Body and the Brain Build an “I”

So far, most of this episode has been argument. Descartes reasoned. Avicenna imagined. Hume looked inward and reported what he found. The Buddhist texts took a person apart with a list.

Over the last forty years, researchers approached the same question from the outside. They stopped asking what the self is and asked something smaller: what has to be working for a person to feel like a person at all? Then they went into the laboratory and disturbed these processes, one by one.

The first thing to say is what they did not find. Neuroscience has not found a single self-centre in the brain — no one place where everything gathers and somebody reads it. Instead, different parts of self-experience seem to depend on interacting systems: systems that track the body, where it is, and what movements should feel like. One way to put the picture together is this: the sense of being one person appears when those systems agree closely enough. Nobody is in charge. There is a great deal of coordination.

So let us go through those layers, from the outside in.

Start with the simplest one. Your body feels like yours.

In 1998, Matthew Botvinick and Jonathan Cohen showed that people could come to feel that a rubber hand lying on the table in front of them was their own hand. In the original experiment, two paintbrushes brushed the hidden real hand and the visible rubber hand in synchrony for ten minutes.

This suggests that the edge of your body — the line where you stop and the world begins — is not as fixed as it feels. When what you see and what you feel are made to disagree, that felt boundary can shift. One common explanation is that the brain works out the boundary from moment to moment, weighing sight against touch — although researchers still debate how much of the illusion this explains.

If you want to go deeper into the evidence, limitations, and references behind this episode, the Companion Paper is available on [LearningPsychiatry.com](https://learningpsychiatry.com).

But owning a hand is only one part of knowing where you are. What about the feeling: I am here, looking out from this position, at this height?

In 2002, Olaf Blanke and colleagues reported a striking case. Electrical stimulation near the temporoparietal junction made a woman feel as if she were above her body, looking down at it. Later reports linked similar out-of-body experiences to disturbances in the same brain region. The details are in the Companion Paper.

These cases did not reveal where the self lives, and a handful of patients cannot tell us how self-location works in everyone. But they show something narrower and still remarkable: the feeling of being located somewhere is not completely fixed. In rare clinical stimulation cases, even a small electrical current was enough to disrupt self-location.

There is another source of self-experience that you cannot see at all: the body as sensed from within.

You do not just see and hear the world. You also sense yourself: your heartbeat, your breathing, your stomach, whether you are too warm, whether something is wrong. This inward sensing is called interoception. Your brain is doing it all the time, whether you pay attention to it or not.

Research from this area suggests that this quiet background sensing may be more than housekeeping. The basic feeling of being here — present, alive, embodied — may partly depend on it. Sensing the body from within may help make experience feel as if it is happening to someone at all.

Body ownership, self-location, and the body sensed from within still leave one familiar question unanswered: who did it? This is the part that gets blamed and praised — the part the man in the boat went looking for in the other boat.

Not “this is my body”, but “I did this.”

Researchers call that the sense of agency: the feeling of being the one who causes things. You might expect it to be obvious. Of course you know when you moved. But the brain seems to work this out too.

One influential explanation is the comparator model. When you move, your brain does not just send a command to the muscles. It also predicts what the movement should feel like, then compares that prediction with what actually arrives from the senses. On this account, a close match contributes to the sense that you caused the movement; a mismatch can weaken that sense.

The everyday version of this is that you cannot tickle yourself. Your brain predicted the touch, so it turned the sensation down. Someone else's fingers were not predicted, so they arrive at full strength.

Put these layers together and you reach a family of self-model theories. Instead of containing a self, the brain may continuously model a body in a place, sensing and acting. Thomas Metzinger puts it simply: the brain's model of you does not feel like a model. It simply feels like you.

That is a proposal, not a proven fact. But it gives us a way to organise the evidence without putting a captain inside the system.

Now return to Avicenna. His floating man was meant to show that self-awareness could remain without bodily sensation. Modern experiments do not disprove a thought experiment, but they do show how deeply ordinary self-experience depends on bodily signals, location, and predicted action. Interfere with those processes, and the ordinary sense of self can bend.

Neuroscience has not found a captain. It has found a very coordinated crew.

Chapter 5 — The Story Called Me

Everything in the last chapter was about the present: where your hand is now, where you are sitting now, what your body is doing now.

But that is not what you mean when you say “I”.

When you say “I”, you also mean a person extended through time: the person who grew up in a particular house, lost someone, was wrong about something for years, changed jobs, changed countries, changed their mind. You are not only here. You have come from somewhere, and you are going somewhere. The word “I” holds all of that together in one syllable.

So this chapter is about the layer that stretches across time. And that layer is not made of stone. It is built from story.

Start with memory, because that is what the story is built from.

We naturally think of memory as a recording. But autobiographical memory is reconstructive: memories of our lives are rebuilt when we need them. Martin Conway and Christopher Pleydell-Pearce described this process as being shaped by our present goals and by what we now believe about ourselves.

Suggestion can also change what people later report as memory. In the classic work by Elizabeth Loftus and Jacqueline Pickrell, some participants later reported a full or partial memory of a childhood event that had never happened. The exact boundary between a rich false memory and a weaker false belief remains debated, but the broader lesson is clear: remembering is not simple playback.

That does not make memory worthless. It means the past reaches us through reconstruction.

One clue to why we tell coherent stories about ourselves came from rare split-brain patients. Michael Gazzaniga observed that when one hemisphere acted on information unavailable to the speaking hemisphere, the person could still give an immediate, confident explanation for the action.

Gazzaniga called this the interpreter: a system that builds a coherent account from whatever information is available. These patients are highly unusual, so they cannot support a general theory on their own. But the idea captures something familiar — humans are very good at explaining themselves, even when the true causes of behaviour are partly hidden.

We do more than remember the past. We also turn memories into a story.

Dan McAdams and Kate McLean call this narrative identity: the story through which a person links past, present, and the future they expect. It has turning points, causes, and a direction. It tells us not only what happened, but what it meant and what kind of person we think we are becoming.

Research suggests that people often feel better when their life story makes sense — especially when they can find some growth or meaning in difficult experiences. But this is only an association. It does not mean that one kind of life story is psychologically “right.”

Depression shows why this matters clinically. People who are depressed often retrieve autobiographical memories less specifically. Asked for a happy memory, they may give a category — “holidays with my family” — rather than one particular afternoon. When specific scenes become harder to reach, a global summary can dominate: “It has always been like this.” The story darkens, and the darkened story can make the future look like more of the same.

Major life events can also disrupt narrative continuity. After illness, bereavement, migration, or divorce, people often say, ‘I am not the same person anymore.’ That sentence makes sense if identity is partly something we continually organise across time.

So the narrative layer is powerful. But it hides an assumption: that everyone naturally experiences life as one continuing story.

The philosopher Galen Strawson rejects that. Some people, he argues, experience life more episodically and do not feel strongly connected to the person they were decades ago. He does not regard that as a deficit.

His challenge matters. If we say a good life must be a well-told story, we turn one way of being a person into a rule for everyone.

So where does that leave us? We now have a body felt from within, a point of view in space, actions that can feel like ours, and a history that gives those moments continuity. None of these is the little person behind the eyes.

The story layer does something especially important: it does not simply describe a finished self. It helps create the sense that there is one continuing author whose past explains the present.

Which takes us back to the river. When the man turned around, he was looking for someone with a reason — a person whose history could explain the collision and make blame possible.

Suppose there had been someone there. He might have said: “I did not see you. I was distracted. The current took me.” The explanation might be immediate, sincere, and coherent.

After this chapter, we have to ask how often we mistake our explanations of ourselves for direct access to the true cause.

Chapter 6 — When the Self Becomes Visible

All the processes we have looked at so far work quietly. You do not feel your brain deciding where your body ends. You do not notice yourself assembling a memory. The layers do their work, they agree with each other, and the result is so smooth that we never see the seams — we just say “I” and get on with the day.

Which is why the seams are easiest to see when something slips.

I want to be careful here, because these experiences are easy to misuse. They do not reveal what the self “really is”. They show something smaller and, for our question, more useful: self-experience is not one single thing. It has parts, and those parts can come apart. Think of them as windows — brief views into a building we normally see only from the outside.

Start with the difference between knowing and feeling.

Some people experience themselves as strange or unreal — as if they were watching their own thoughts and actions from a small distance rather than living them. This is called depersonalisation. People may describe it as: “Everything is behind glass. I am going through the motions. My hands are doing this, and I am somewhere else, watching.” And here is the part that matters. They still know what is real. Ask them directly, and they will tell you that they are real, that they are here, that the hands are theirs. They know that the feeling of unreality is a feeling, even though it can be overwhelming. What is missing is not the knowledge. It is the feeling that normally goes with it.

So knowing that you exist and feeling real can come apart. Most of us rarely notice that, because in ordinary experience they usually arrive together.

Now take the second window. It leads straight back to Chapter 1.

In Chapter 1, we separated two feelings: that a thought is mine, and that I am the one producing it. In everyday life, they usually arrive together. Sometimes a thought is clearly present in a person's own mind but does not feel as if they produced it. Nobody else can hear it. The person knows it is happening in their own mind. But it feels placed there. This experience is called thought insertion, and one influential way of understanding it is that ownership and agency have come apart.

This is hard to imagine from the outside. It turns a distinction philosophers have debated for centuries into a lived experience: ownership without agency, on this reading — although not everyone agrees that this is the cleanest way to describe it. The thought is here, in my mind, and I did not make it.

The question from the river has turned inward: the thought is here, but the sense of an agent behind it is missing.

The third window is about identity over time and how other people recognise us.

In 2015, Nina Strohminger and Shaun Nichols surveyed relatives of people with different brain diseases. The largest perceived changes in identity were linked less to memory loss itself than to changes in moral traits such as kindness and honesty.

It is one study, and it measures relatives' judgements rather than the patient's inner experience. But it suggests that when we decide whether someone is 'still themselves', character can matter more than memory.

Taken together, these windows show that the seamless "I" has distinguishable dimensions: knowing and feeling real, owning and generating a thought, remembering a life, and being recognised by others as the same person. They usually arrive together. Sometimes they do not.

The smoothness of ordinary selfhood does not mean that nothing is happening. It is what coordinated processes feel like when they work together.

Ending — No Little Person, Not Nothing

Let us go back to the river one last time.

We have been watching the empty boat. Now watch the man. He turns around ready to shout, finds nobody there, and something lets go. The pain and the cold water remain. What disappears is the picture of a chooser behind the collision.

The idea explored in this episode is that when we look for a comparable chooser inside ourselves, we do not find a little person either.

But we do not find nothing. We find a body sensed from within, a point of view, actions that can feel like ours, memories reconstructed as we reach for them, and stories that connect one moment to the next. Self-model theories also describe continuously updated representations of body, position, and action. What we do not find is an extra captain in addition to all of that.

The chariot gives us a simple way to think about the self. A chariot is real, even though there is no extra "chariot" hidden somewhere inside its parts. The same may be true of us.

There may be no little person behind your eyes. But that does not mean nobody is there. There is a real person — with a body, memories, relationships, habits, a history, and a life that continues to change.

So what holds all these parts together? One answer is that the brain is constantly linking information about the body, our actions, our memories, and the world around us. When these processes work together, they create the continuous experience of being one person — "me". We understand some of this, but not the whole picture yet.

And that leaves one final question. If there is no separate captain inside, then what are we really blaming when we blame ourselves? And what are we forgiving when we forgive someone else?

Maybe the self is not the captain of the boat. Maybe it is the whole journey — the boat, the river, the weather, the memories of earlier crossings — all coming together in the strange experience we call "I".

I'm Said Fard, a psychiatrist here in Sweden. Keep learning.