

“I Became a Neuroscientist to Understand Consciousness. I Spend My Days Writing Grant Proposals.”

Piece #009 · TESTIMONY · The Infinite Puzzle

What follows is a composite account — assembled from conversations with researchers at the edge of the inner/outer divide. The voice is one. The experience is many.

I can tell you the exact moment I decided to become a neuroscientist.

I was nineteen, reading in my dorm room, and I came across a sentence in a philosophy of mind textbook that stopped me entirely. The sentence described what philosophers call the hard problem of consciousness — the question of why physical processes in the brain give rise to subjective experience at all. Why there is something it is like to be you, reading this, right now.

I read the sentence three times. Then I closed the book and sat with the feeling of it.

The feeling was not confusion. It was recognition. The sense that this — *this specific question* — was the most important question anyone could ask, and that I had just found the field that asked it.

I was wrong about the field. Not wrong about the question.

What the Field Actually Is

Neuroscience, as it is practiced inside major research institutions, is not primarily the study of consciousness. It is the study of measurable neural correlates of behavior and cognition — an extraordinarily difficult and important scientific project, relentlessly rigorous, deeply specialized, and organized almost entirely around questions that can be operationalized in a way that generates publishable findings.

This is not a criticism. It is a description of how science necessarily works. You can only study what you can measure. You can only publish what you can defend. You can only secure funding for what you can promise, in advance, to deliver.

The question of consciousness — of why physical processes give rise to felt experience, of what the structure of inner experience actually is, of how subjective states relate to objective measurement — is not a question that fits comfortably inside these constraints. It is too large. Too foundational. Too resistant to the kind of experimental design that produces a clean p-value and a clear contribution to the literature.

And so what happens, to almost everyone who enters the field with that question burning, is a slow, polite displacement. You choose a manageable sub-question. You build an experimental paradigm. You publish. You apply for the next grant. You supervise students who are doing the same.

The original question doesn't go away. It migrates — into the margins of your thinking, into the papers you read on weekends, into the conversations you have at conferences with the three other people who are still asking it in any serious way. It becomes a private obsession that your CV doesn't quite capture.

The question I actually want to ask doesn't fit in a grant proposal. It doesn't fit in a paper. It barely fits in a conversation.

The Specific Shape of the Frustration

I want to be precise about this, because the frustration is not what people outside academia usually imagine.

It is not that the institution is hostile to deep questions. Many of the people inside it are brilliant and genuinely motivated by the largest possible versions of the questions they work on. The frustration is structural, not personal.

The structure produces a specific kind of pressure: toward questions that are already partially answered. The closer you are to existing literature, the easier it is to frame a contribution, defend a design, and place a paper. The further you move from the existing literature — toward genuinely new questions — the harder every step becomes. Reviewers don't have a framework for evaluating it. Funding bodies can't assess the risk. The journal that would publish it doesn't exist yet.

The questions at the frontier of the inner/outer problem — What is the structure of felt experience? How can inner-state signal be captured in a form that outer-world systems can receive? What would it mean for an AI to have a model of the experience it is trying to serve? — are not at the margin of existing literature. They are off the edge of the map.

Doing science off the edge of the map, inside an institution designed for the map, is a particular kind of exhaustion. It requires translating everything you're thinking into terms that fit a framework you've already half-outgrown.

The Conversations That Don't Happen

There are conversations that happen at conferences in the hallways and at late dinners, after the formal sessions have ended and the language becomes less guarded.

In those conversations, the question changes. It stops being “what is the neural correlate of x?” and becomes something more like: “do you actually think we're going to understand

consciousness from inside this paradigm, or are we building an increasingly detailed map of a coastline we're not actually trying to cross?"

I have had versions of this conversation with people at every level of the field — graduate students who are beginning to sense the gap, tenured faculty who have been living inside it for decades, and a few people who have left academia entirely and are trying to build something in the space that the institution cannot hold.

The most honest version of the conversation usually surfaces the same recognition: the work the field is doing is real and valuable and insufficient. It is insufficient not because the scientists aren't good enough, but because the questions that would complete the picture are questions the institution is not structured to ask.

What would it look like to build the instruments for inner-state translation outside the institution? What would it look like to take the theoretical frameworks that the field has spent decades developing — interoception, predictive processing, active inference, constructed emotion — and actually apply them to the problem of making inner experience legible to outer-world systems?

Nobody in the institution has time to answer that question. It doesn't fit in the grant cycle.

The Moment I Started Looking Elsewhere

Last year I was reviewing a funding proposal from a colleague — a genuinely promising project studying the physiological correlates of decision-making under uncertainty. The science was solid. The design was careful. The potential contribution to the literature was clear.

And I realized, halfway through the review, that I didn't care.

Not that I didn't respect the work. I did. But I had reached the point where my honest response to a well-designed study within the existing framework was a kind of neutral appreciation rather than excitement. The question it was answering was real. It was not the question I thought mattered most.

That gap — between the work available to me inside the institution and the work I believed was most needed — had been widening for years. The review was just the moment it became impossible to paper over.

I started reading differently after that. Less toward the center of the field, more toward its edges. I followed threads into somatic psychology, into contemplative neuroscience, into the applied philosophy of mind, into the small and loosely connected community of people who were trying to build instruments for the inner world rather than just map it from the outside.

I started asking a different question: not "what can I measure from here?" but "what needs to be built, and who is building it?"

What Needs to Exist

I am still inside the institution. I am not sure yet whether I will stay.

But I know, with a clarity that the institution itself cannot receive, what needs to exist that does not.

It is not a better study. It is not a more sophisticated neuroimaging protocol or a larger sample or a more rigorous experimental design. Those things are valuable. They are not the gap.

The gap is a translation layer — between the inner-state data that the neuroscience is now sophisticated enough to generate frameworks for, and the outer-world systems that could act on that data if they had any way to receive it. The gap is between the map and the road. Between the territory and the infrastructure.

Nobody in the institution has a grant for that. The application form doesn't have a category for it. The reviewers don't have a rubric.

But the need is real. The tools are close enough to touch. And somewhere outside the institution — at the edge where builders and researchers and people who refuse to choose between rigor and ambition are beginning to find each other — the work is starting.

I can feel that it is starting. Which, for a neuroscientist who studies inner experience, is exactly the kind of signal I have learned to trust.

The Infinite Puzzle is the public record of the PRIMORDIAL Project — building the translation infrastructure between the inner and outer worlds. If a piece of this puzzle belongs to you, the door is open.

[← Piece #008] | [Subscribe] | [Piece #010 →]