

# Is a T-cell Any More Real Than a T-test? A Plea for Epistemic Openness

Ravi Varadhan, Johns Hopkins University

During a journal club discussion, an oncology colleague claimed cytotoxic T-cells were "the real deal" while dismissing a T-test result as "just math, something we made up." This crystallizes a presumed divide in biomedical culture: biological entities are real, statistical tools are abstract. This essay examines that distinction. We argue that both the T-cell and the T-test are powerful abstractions built from layers of observation, inference, and interpretation. The T-cell, far from being a brute fact of nature, is an operational definition: a cell that expresses specific surface proteins like CD3, identified not by sight but by affinity for fluorescent antibodies. The T-test translates noisy biological measurements into probabilistic inferences about treatment effects that would otherwise remain invisible.

Drawing on Hacking's distinction between entity realism and theory realism, we propose that scientific concepts exist on a continuum of abstraction. A concept's reality is best judged by its usefulness—its capacity to help us describe, predict, and intervene within the prevailing paradigm. Both the T-cell and the T-test pass this test.

Yet we part ways with Hacking's ontological confidence. What matters is not whether these concepts are "really real" in some paradigm-independent sense, but that they are paradigmatically indispensable—load-bearing abstractions without which modern biomedicine would be substantially weakened. This position, a kind of pragmatism chastened by history, draws on but also diverges from constructive empiricism, structural realism, and the natural ontological attitude.

From this view follows a practical imperative: epistemic openness. If all scientific concepts are paradigm-dependent, and if paradigms shift, then dogmatism about which abstractions are "real" is indefensible. Recognizing both as abstractions may reframe interdisciplinary collaboration and temper the reflex to privilege whatever feels more concrete. Paradigmatic indispensability—not ontological certainty—is all any scientific concept needs to claim.

**Keywords:** Philosophy of science, Scientific realism, Biomedical epistemology, Statistical inference, Scientific paradigms, Paradigmatic indispensability

## Introduction

During a recent journal club, my oncology colleague Dr. James pointed at a flow cytometry plot and said: "See those? The cytotoxic T-cells. The real deal. These are what's going to save the patients." Then he turned to me. "And your T-test? It's a useful tool. I'll give you that. But it's just math. Something we made up."

In that moment, he captured a profound divide I've struggled with throughout my career: the presumed gulf between the "real" objects of biology and the "abstract" tools of statistics. This interaction forces us to confront a fundamental question: Is the conceptual model we call a "T-cell" inherently more "real"

than the conceptual model we call a "T-test"<sup>1</sup>? This essay is an attempt to unsettle that distinction; to examine whether the boundary between biological reality and mathematical abstraction is as firm as is commonly taken for granted.

### **The T-cell is an Abstraction Too**

The T-cell is undeniably a cornerstone of modern immunology<sup>2</sup>, a concept validated by remarkable therapeutic successes like CAR-T. However, its apparent self-evidence dissolves under closer inspection. Consider how we actually identify a T-cell in the laboratory. We cannot simply "see" one. What we call a T-cell is, in practice, an operational definition: a cell that expresses specific surface proteins like CD3, with subsets distinguished by proteins such as CD4 or CD8, which we identify not by sight, but by their affinity for fluorescent antibodies. The cell is defined by what sticks to it. Operationally, a 'T-cell' is a gated CD3+ event whose identity depends on instrument settings, compensation algorithms, and analyst decisions. Modern standardization mitigates subjectivity but does not abolish it altogether, underscoring that even our 'real' cells are seen through layers of instrument and inference.

Let's trace this back further. The "cytotoxic T-cell" is a functional concept; a cell that kills infected or malignant cells. Its laboratory identification is a CD3+CD8+ *event* on a flow cytometer. But what is the *raw* experience underlying that identification? A series of voltage pulses generated when a laser hits a cell, scattered and captured by detectors, then digitized into a coordinate on a scatterplot. What my colleague calls a T-cell is, in the moment of observation, a dot on a graph.

The irony, of course, is that my colleague was not looking at a cell during journal club. He was looking at a *statistical visualization* of light scatter and fluorescence, which is a graph constructed by software that applies compensation algorithms, gates, and probability calculations. His T-cell was, in that moment, as much an inference from data as any p-value I might calculate from a T-test. The difference is one of familiarity, not fundamentality.

The T-cell is therefore not a brute fact of nature, handed to us directly by our senses. It is a powerful and coherent model that integrates decades of experimental data from microscopy through flow cytometry to gene expression profiling into a single, useful concept. It is an abstraction so successful that we have forgotten it is one.

### **The T-test is an Abstraction Too**

---

<sup>1</sup> The t-test was developed by William Sealy Gosset (1876–1937), an English statistician working for Guinness Brewery in Dublin. Guinness prohibited employees from publishing under their own names to protect trade secrets, so Gosset used the pseudonym "Student." The test has been formally known as Student's t-test. Throughout this essay, I capitalize the "t" in t-test for visual consonance with "T-cell."

<sup>2</sup> I use the T-cell as a metonym for the broader family of biological constructs that power modern biomedicine such as receptors, cytokines, signaling pathways, and gene regulatory networks. Each is operationally defined, instrumentally detected, and layered with inference. The point is not the specific entity but the way biomedical culture treats such constructs as self-evidently real.

The T-test, in contrast, is patently abstract.<sup>3</sup> It is a formula, a calculation, a set of assumptions involving normality and variance. Where the T-cell presents itself as a thing we discover, the T-test presents itself as a tool we invented. My colleague's critique that "It's just math. Something we invented" seems, on its face, entirely fair. But let us examine what the T-test actually does.

Consider a simple experiment: twenty mice with tumors, ten treated with a new immunotherapy, ten given a placebo. We measure tumor volumes. The treated mice have smaller tumors on average. But biology is messy. Some treated mice have larger tumors than some untreated mice. The signal is buried in noise. Is the difference real, or just random variation?

This is where the T-test enters. It is a powerful mathematical abstraction that allows us to see differences between two sets of biological measures in the presence of noise. It cannot tell us with certainty whether the treatment works—no tool can—but it can quantify the probability that an observed difference this large would appear if the treatment actually did nothing. It cuts through the noise and helps us decide whether to believe our eyes. The T-test begins with concrete measurements: the tumor volumes themselves. These numbers are not abstractions; they are physical readouts from calipers or imaging systems. The test processes these numbers through a formula, producing a *t*-statistic and, from that, a *p*-value. That *p*-value is not a comment about the numbers themselves, but an inference about the process that generated them. It is our best attempt, given limited data, to distinguish genuine biological effect from random fluctuation. Dr. James might object that this is where statistics overreaches. Descriptive statistics such as plots, tables, and summaries stay close to the data and feel trustworthy. Inferential statistics leap beyond the observed sample to make claims about the unobserved world. But this leap is not away from reality; it is toward the reality we actually care about. The twenty mice in front of us are not the point. The point is the next thousand patients we will never directly observe. Description tells us what happened here; inference tells us what to expect out there. It is the very thing that makes statistics useful for medicine.

Notice what has happened here. Just as the flow cytometer translates laser flashes into a scatterplot, the T-test translates noisy experimental measurements into a probability. Just as my colleague looks at a graph and infers the presence of T-cells, I look at a *p*-value and infer the presence of a treatment effect. Both of us are staring at numbers on a page. Both of us are making inferences about an unseen reality. Both tools, the flow cytometer and the T-test, help us see what would otherwise be invisible. The key difference is that the T-test's abstractions are more obviously mathematical. But is that the same as being less real? The *p*-value from a well-conducted trial has guided countless decisions to move drugs forward or abandon dead ends. It has predicted patient outcomes. It has, indirectly, saved lives. If a concept's reality is measured by its power to intervene in the world, the T-test stands on remarkably solid ground.

---

<sup>3</sup> The T-test is used as a metonym for the broader family of statistical inference techniques that power contemporary biomedicine such as Bayesian hierarchical models, generalized linear and mixed models, survival models, and even the recent machine-learning approaches used for causal or predictive purposes. The point is not the specific test but the practice of principled inference from noisy data.

The T-test is therefore not a mere game of numbers, disconnected from biological reality. It is a tool for detecting patterns that matter, patterns of treatment effect, patterns of harm, patterns of benefit. It is an abstraction, of course. But it is an abstraction about the same noisy, complex world that the oncologist inhabits, and it earns its keep through its utility.

### **The Continuum of Reification and the Primacy of Usefulness**

Scientific concepts exist on a continuum of abstraction, and it is worth tracing where our two protagonists, the T-cell and the T-test, sit along it. At the most immediate level is direct perception: the palpable tumor, the visible rash, things we encounter through our senses without mediation. One step removed is the instrumental level, where optics intervene between us and the object, e.g., a bacterium under a microscope, a cell visible only through a lens. Further still is the operational or functional level, where we no longer see the thing directly but infer it from assay readouts: a color change, a dot on a scatterplot. Beyond that lies the inferential level, where we compare groups and extract meaning from the difference, e.g., a treatment effect, a signal of toxicity, or a p-value. And at the furthest reach is the theoretical level, where concepts like immune memory or natural selection are inferred not from any single observation but from patterns of behavior accumulated over time.

What is striking is that the T-cell does not sit at a single point on this continuum. We can see it through a microscope, placing it at the instrumental level. We define it by its surface markers, placing it at the operational level. We theorize about its role in immunological memory, placing it at the theoretical level. The T-test, meanwhile, operates primarily at the inferential level but it draws its inputs from the operational and instrumental levels, and its outputs inform theoretical conclusions. Neither concept belongs exclusively to the "real" or the "abstract" end of the spectrum. They are different points on the same continuum, and the continuum itself is what science navigates every day.

The philosopher Ian Hacking (1983) drew a useful distinction between two separate questions we can ask about scientific knowledge: whether our theories are true, and whether the entities those theories invoke actually exist. These are independent questions. One can believe that electrons exist because we can manipulate them, build devices that depend on them, predict their behavior, without committing to any single theory of electron physics being definitively true. It is Hacking's second question, about entities rather than theories, that concerns us here. We are not asking whether immunological theory or frequentist statistics is true in some final sense. We are asking something more modest and more tractable: whether the T-cell and the T-test are real in the sense that matters, which is whether they reliably help us understand, predict, and intervene in the world. Following Hacking's distinction between 'representing' and 'intervening,' one way to cash out the reality of a scientific model is by what we can reliably do with it. We believe that cytotoxic T-cells are real not merely because we photograph or gate them, but because we intervene on them, expand them, engineer them (CAR-T), and observe reproducible alterations in tumor burden as a result. That is entity realism: practical, intervention-grounded commitment. By the same token, while a T-test is not a real entity, its entrenched use is justified by a parallel form of interventionist success: it reliably guides action. When randomized trials

cross prespecified thresholds, clinicians change practice; mortality falls and harm is averted. The map is not the territory, but it is a good map when it repeatedly gets us to where we intend to go.

A scientific concept's position on this continuum, then, does not determine its reality. What determines reality is its usefulness. A concept is real in the only way that matters for scientific practice - if it reliably helps us describe, predict, and intervene in the world. The T-cell passes this test. It allows us to predict that a patient with low CD4 counts will be vulnerable to infection. It allows us to intervene by engineering CAR-T cells to attack tumors. These are not trivial achievements. They are why my colleague calls the T-cell "the real deal." But the T-test passes the same test. It allows us to predict, with quantified uncertainty, whether a new drug is likely to benefit patients. It allows us to intervene by stopping a trial early when the evidence for benefit or harm becomes clear. Clinical guidelines, regulatory approvals, and treatment decisions rest on the foundation of statistical abstractions like the T-test. Countless patients are alive today because someone, somewhere, trusted a T-test enough to change practice.

The two abstractions are not rivals: the T-cell gives us a biological mechanism to target, while the T-test gives us confidence that targeting it actually works. One without the other is either blind faith or impotent observation. Together, they form the basis of evidence-based medicine. Both help us "see" things we cannot perceive with our naked senses: we see them with our mind's eye. The oncologist looks at a scatterplot and sees a cytotoxic T-cell hunting cancer. The statistician looks at a p-value and sees a treatment effect that would be invisible through casual observation. Both are engaged in the same essential act: using abstraction to extend the reach of human perception. But why does Dr. James's intuition feel so natural? The answer lies not in logic alone, but in history.

### **A Conversation Across Centuries**

The divide my colleague expressed so casually is not merely a personal opinion. It reflects a fault line running through the heart of modern biomedicine. Officially, biomedicine is a unified paradigm—a synthesis of biological mechanism and statistical inference—that allows us to understand, predict, and intervene in disease. But this unity is a marriage of convenience, not a fusion of equals. Two distinct epistemic traditions cohabit uneasily within it: one rooted in the laboratory, the other in mathematics; one that trusts what can be seen under a microscope, the other that trusts what can be calculated from noisy data; one that asks 'what is the mechanism?' and the other that asks 'what is the evidence?'

The philosopher Thomas Kuhn (1962) argued that science does not progress through the steady accumulation of neutral facts, but through the dominance of paradigms: shared frameworks of assumptions, methods, and values that determine what counts as a legitimate question, a valid method, and a real entity. Within biomedicine, the dominant tradition has always been biological. The natural unit of explanation is the cell, the receptor, the molecule. Statistics entered biomedicine in the late nineteenth century, emerging from the work of Galton, Pearson, and (later) Fisher on heredity and biological variation. It proved indispensable but it never fully belonged. It remained a visitor in someone else's house: useful, tolerated, but not quite family.

My colleague's intuition that T-cells are real while T-tests are merely useful is not a philosophical conclusion he reasoned his way to. It is the inherited bias of the biological tradition, so deeply embedded that it feels like common sense. Kuhn would recognize this immediately: within normal science, foundational commitments are invisible precisely because they are never questioned. The journal club exchange between an oncologist and a statistician is a glimpse of what happens when two epistemic traditions—long cohabiting, never fully reconciled—collide over a simple question: what is real?"

This question of what is real has been asked for over two millennia. It started with Plato. When my colleague points at a flow plot and declares T-cells "the real deal," he is making an implicit claim about where reality resides: in the physical, the tangible, the biological. Plato would have found this bemusing, for it is precisely backwards. For Plato, the physical world is a shadow play, a flickering cave of imperfect copies. True reality, eternal, unchanging, perfect, resides in the abstract forms. A circle I draw on a whiteboard is a flawed approximation; the *idea* of a circle is what is real. A particular T-cell in a particular patient will die, mutate, be destroyed. The *concept* of a T-cell, the functional role it plays in our immunological models, has a kind of permanence no individual cell can match. From a Platonic perspective, my colleague has it exactly backwards: the mathematical and conceptual structures are *more* real, not less. Legend has it that above the door of Plato's Academy was inscribed: "Let no man ignorant of geometry enter here." For Plato, geometry was not merely useful; it was a preparation for encountering reality itself. To think clearly about the world, one must first master the abstractions that reveal its true structure. My colleague, who walks daily through the doors of modern medicine, might not realize how much of what he sees depends on ancestors who honored that inscription.

Galileo, standing at the birth of modern science, offered a more practical reconciliation. He wrote that the great book of the universe "is written in the language of mathematics, and its characters are triangles, circles, and other geometric figures, without which it is humanly impossible to understand a single word of it." For Galileo, mathematics was not an escape from reality but the very means by which reality becomes intelligible. Nature keeps her truths hidden; we cannot read them unless we learn her language. My colleague sees a T-cell attacking a tumor. Galileo would ask: How do you *know* it is attacking? Because you have measured killing rates, compared them to controls, calculated differences that exceed what chance would predict. You have read nature's book. And the book is written in mathematics.

In 1960, the physicist Eugene Wigner (1960) published an essay that gave the tradition a name: "The Unreasonable Effectiveness of Mathematics in the Natural Sciences." Wigner marveled that, mathematical concepts, often developed for pure aesthetic reasons with no application in mind, repeatedly turn out to be exactly what physicists need to describe the universe. Why should abstract algebra, invented by mathematicians playing with symbols, perfectly capture the behavior of subatomic particles? Why should non-Euclidean geometry, dreamed up as a logical curiosity, become the very fabric of spacetime in general relativity? And why, we might add, should a statistical test developed by a brewer in Dublin to assess the quality of stout prove so effective at evaluating cancer treatments (Student 1908)? Wigner called this a "miracle" something we cannot explain but cannot deny. Einstein, too, wondered at it: "How can it be that mathematics, being after all a product of human thought which is independent of experience, is so admirably appropriate to the objects of reality?"

Notice what these thinkers share. None of them denies the reality of the physical world. But all of them insist that mathematical abstraction is not a veil between us and reality, but a magnifying lens. It allows us to see patterns, structures, and relationships that are simply invisible to the naked senses. The T-cell is real. But so is the difference between treated and untreated groups that only a T-test can reveal. That difference is not a number on a page; it is a fact about the world. The T-test simply lets us see it.

My colleague says we statisticians "just made up" our tools. In one sense, he is right. We did make them up, just as biologists made up the concept of a T-cell, and geometers made up the concept of a triangle, and Galileo made up the language in which he read the book of nature. But "making up" is too simple. These creations are not arbitrary. They are shaped by something outside themselves. A model of T-cell that failed to predict immune behavior would be discarded. A statistical test that consistently led to wrong conclusions would not survive. The fact that our abstractions work - that they explain, predict, and heal suggests they are not mere fictions. They are more like keys: we shape them, but they must fit the locks they encounter.

Wigner called this the "unreasonable effectiveness" of mathematics. Perhaps it is unreasonable. Perhaps it is a miracle. Or perhaps, as Plato might have said, there is an order to reality that our minds are somehow tuned to perceive. We do not invent triangles; we discover their properties. We do not invent logic; we uncover its necessary truths. The same may hold for the concepts of biology and statistics. We build them, yes. But they work because they correspond, however imperfectly, to something real. That, perhaps, is the deepest source of wonder: not that our inventions randomly prove useful, but that the universe appears to be structured in a way that makes it intelligible to minds like ours.

### **Seeing with the Mind's Eye**

My colleague Dr. James pointed at a flow cytometry plot and saw something real: cytotoxic T-cells, poised to save his patient. He turned to me and saw something abstract: a T-test, a human invention, useful but not quite real. I have spent this essay trying to unsettle that distinction.

The T-cell, for all its therapeutic power, is not a brute fact handed directly to our senses. It is an operational definition, a functional concept, an inference from fluorescent signals and voltage pulses rendered as a dot on a graph. It is an abstraction so successful that we have forgotten it is one. My colleague sees with his mind's eye a cell attacking cancer; what his physical eye sees is a scatterplot. The T-test, for all its mathematical formality, is not a mere game of numbers. It is a tool that allows us to see differences buried in biological noise, to detect patterns that would otherwise remain invisible. It is an abstraction that extends our perception, just as a microscope does. The p-value from a well-conducted trial has guided treatments, saved lives, and shaped the very standards of evidence my colleague relies upon when he calls a therapy "proven."

Plato claimed that mathematical objects are more real than sensory ones. Galileo insisted that nature's book is written in mathematics. Wigner and Einstein marveled at the unreasonable effectiveness of this language we invented, which somehow unlocks the universe's deepest secrets. Each, in his own way, understood that the mind's eye sees more clearly and farther than the physical one.

Consider what we have achieved by weaving together abstractions like the T-cell and the T-test. The T-cell names a biological reality we cannot directly see: a surveillance system that patrols our bodies, distinguishing self from foreign, remembering enemies it encountered decades ago. The T-test names a statistical reality we cannot directly see: a signal of treatment effect rising above the noise of biological variation. Together, they exemplify something remarkable about science itself. We are creatures of limited senses, trapped in a narrow slice of the electromagnetic spectrum, perceiving only a fraction of what is. And yet, by building abstractions, biological and mathematical, we have extended our perception across scales of magnitude and complexity. We see cells. We see patterns amidst randomness. We see treatment effects. We see the interior of stars and the echo of the Big Bang. We see with our mind's eye what no instrument could ever see.

This is not a matter of biology being "real" and statistics being "abstract." Both are abstractions. Both are real. And both are tools in a shared endeavor that is, when you step back and look, almost unbelievable: that beings like us, born into a universe we did not make, can build models that unlock its secrets, heal its diseases, and fill us with wonder at its intelligibility. Dr. James and I are not engaged in different enterprises, one dealing with the real and the other with the abstract. We are engaged in the same enterprise: using our mind's eye to perceive what our naked senses cannot. He sees T-cells in a scatterplot. I see treatment effects in a T-test. Both of us are reading the book of nature using the language we have mastered.

### **Beyond Entity Realism: Paradigmatic Indispensability**

In the section "The Continuum of Reification" I leaned on Hacking's distinction between entity realism and theory realism to argue that both T-cell and T-test though they may lie on different points along the continuum are *real* because of their usefulness. This allowed us to sidestep the unanswerable question of whether our theories are true in some final sense, and to focus instead on whether the entities they invoke earn their place through practical success. The T-cell earns its place because we can manipulate it; the T-test earns its place because it reliably guides action. So far, so good. But here I part ways with Hacking. His argument, at its strongest, is an *ontological* claim: entities we can manipulate are *real*, period, independent of our theories about them. I do not want to go that far. My claim is more modest, and perhaps more Kuhnian: I do not know whether T-cells or T-tests are "ontologically real" in some paradigm-independent sense. I am not sure the question is answerable, or even well-formed.

What I *do* claim is this: both the T-cell and the T-test are epistemically indispensable within our current paradigm. They are not optional tools we could discard without loss. They are load-bearing abstractions that make modern biomedicine possible. The T-cell allows us to understand, predict, and intervene in immune function. The T-test allows us to extract signal from noise and make decisions under uncertainty. Neither can be eliminated without the edifice collapsing. This is a pragmatist position, in the tradition of William James and John Dewey: the meaning of a concept lies in its practical consequences, and its "truth" is a matter of its usefulness in helping us navigate the world (James 1922). It is also consonant with van Fraassen's constructive empiricism (van Fraassen 1976), which holds that science aims not at truth but at empirical adequacy—theories that "save the phenomena" (predict/explain

observations adequately) without requiring belief in the reality of unobservable entities. And it overlaps with structural realism, which suggests that what persists across paradigm shifts is not the entities themselves but the mathematical structures that relate them (Worrall 1989). Finally, it shares something with Arthur Fine's 'Natural Ontological Attitude' (Fine 1984), which counsels that we stop asking metaphysical questions about whether scientific entities are 'really' real and simply accept science's claims at face value. But where Fine recommends quietism, I counsel historical awareness: paradigmatic indispensability is adequate for now, but history teaches us it may not be permanent. In this sense, I am probably staking out a position of my own. Let me elaborate on the similarities and distinctions, addressing each philosopher in turn.

What I have called *paradigmatic indispensability* draws on this philosophical tradition but makes a distinct claim: that being load-bearing within a working paradigm is all any scientific concept can claim, and all it needs to claim. With Hacking, I share the conviction that intervention matters—we trust what we can manipulate and what reliably guides action. But I stop short of his ontological conclusion. Hacking wants to say: because we can spray electrons and observe their effects, electrons are *real*, full stop. I am not convinced the question admits of such a settled answer. What we *can* say is that electrons, like T-cells and T-tests, are indispensable within the current paradigm. That is an epistemic claim about their function, not an ontological claim about their essence. Hacking asks "Is it real?" and answers with an interventionist criterion. I ask a different question: "Can we do without it?" The answer—*today*—is no.

With van Fraassen, I share a certain epistemic humility. The aim of science need not be Truth with a capital T; empirical adequacy—saving the phenomena—is a perfectly respectable goal. But constructive empiricism asks a question I am not asking: "What should we believe?" Van Fraassen answers: only what is observable. Paradigmatic indispensability asks a different question: "What can we not do without?" The answer is not about belief but about function. T-cells and T-tests are load-bearing concepts in the current scientific edifice. Remove them and the structure collapses. That is a claim about practice, not about ontology or even epistemology in the traditional sense. Where van Fraassen counsels suspension of belief in unobservables, paradigmatic indispensability counsels something more like *temporal humility*: these concepts work now, but history whispers that they may not work forever. The stance is not agnosticism but *alertness*—a willingness to see today's indispensables as possibly tomorrow's curiosities.

With Worrall, I share the intuition that *something* persists across theory change. Structural realism suggests that what survives paradigm shifts is mathematical structure: Fresnel's equations, for example, remain valid even as the ontology of the luminiferous ether dissolves. I find this plausible but perhaps too narrow. Sometimes what persists is not mathematical structure but functional role, causal patterns, or a way of seeing. The T-cell concept may one day be superseded by a systems-level understanding that renders discrete cell types obsolete—but whatever replaces it will have to account for the same phenomena: immune surveillance, memory, cytotoxicity. The continuity may be at the level of *what the concept does* rather than *how it is mathematically expressed*. This is a weaker claim than structural realism makes, but perhaps a safer one.

With Fine, I share the desire to stop torturing ourselves with unanswerable metaphysical questions. Fine's "natural ontological attitude" counsels us to accept science's claims at face value and move on. There is something appealingly sensible about this. But I find it too quietist. History shows that science's claims change, sometimes dramatically. The natural attitude, for me, is not passive acceptance but *engaged provisionality*: we rely on these concepts, we build our careers and our patients' lives on them, and we also remain alert to the possibility that they may one day be superseded. That alertness is not skepticism; it is intellectual maturity. Fine says "stop asking." I say "keep asking, but don't let the asking paralyze you."

Paradigmatic indispensability, then, is a kind of pragmatism chastened by history. It asks not "Is it real?" but "Does it work?" and adds, "And for how long has that been true? And how long might it continue?" It draws on Hacking's attention to intervention, van Fraassen's humility about unobservables, Worrall's sense that something persists, and Fine's impatience with metaphysics—but it synthesizes them into something none of them, individually, quite says. It is a modest position, but I think it is all any scientific concept can claim, and all it needs to claim.

Paradigmatic indispensability implies epistemic openness. If all scientific concepts are load-bearing abstractions within a paradigm, and if the history of science teaches that paradigms shift, then dogmatism about which abstractions are "really real" becomes intellectually indefensible. The philosophical position leads naturally to a practical stance: we must hold our concepts, whether biological or statistical, with sufficient humility to recognize their contingency, while taking them seriously enough to build careers and save lives upon them. This is not skepticism; it is the mature recognition that today's indispensables may be tomorrow's curiosities, and that the best we can do is work faithfully within our paradigm while remaining alert to the shifts that may occur.

At this point, a skeptic might ask: if our criterion is paradigmatic indispensability—being useful within a working paradigm—what happens when paradigms shift? Will "T-cells" seem as quaint to future immunologists as "animal spirits" seem to us? History of science suggests we should take this question seriously. It offers humbling precedents (Laudan 1981). Consider phlogiston, the substance eighteenth-century chemists believed was released during combustion. Phlogiston was not idle speculation; it was paradigmatically indispensable. It explained why metals transformed when burned, why charcoal was rich in combustible potential, why a candle extinguished in a sealed jar. Chemists manipulated phlogiston-rich and phlogiston-poor substances, made predictions, designed experiments. Phlogiston was, by Hacking's interventionist criterion, as real as anything in chemistry. Priestley, one of the greatest chemists of his age, went to his grave defending it. Then Lavoisier introduced oxygen, and phlogiston vanished not because the observations changed, but because a better framework made it unnecessary.

Or consider caloric, the invisible fluid thought to carry heat from hot bodies to cold ones. Caloric explained conduction, radiation, and the expansion of gases. It was quantifiable: engineers calculated caloric flow when designing steam engines; Carnot's foundational work on thermodynamics was conducted entirely within the caloric framework. It was useful; it was predictive; it guided intervention.

Yet caloric dissolved when thermodynamics reframed heat as molecular motion rather than a substance. The same observations that once demanded caloric now demanded kinetic energy.

Even in medicine, the four humors: blood, phlegm, yellow bile, and black bile were indispensable for two millennia. Hippocrates systematized them; Galen built an entire physiological edifice upon them. Physicians balanced humors, prescribed bloodletting, adjusted diets according to humoral theory. The framework explained temperament, disease, and cure. It was coherent, teachable, and actionable. Generations of patients were treated on the basis of humoral reasoning. It was also, we now believe, entirely wrong about the entities it invoked. Yet within its paradigm, it was as real as any modern biomarker. A Renaissance physician who dismissed the humors would have been as professionally suspect as a modern oncologist who dismissed T-cells.

These are not examples of bad science. They are examples of *normal* science, conducted rigorously within the frameworks available at the time. The practitioners of phlogiston chemistry and humoral medicine were not fools; they were experts operating at the frontier of knowledge. Their concepts were indispensable - until they weren't.

I am not predicting that T-cells or T-tests will be abandoned. But epistemic humility requires us to hold even our most successful concepts with open hands. As we have just seen, the history of science teaches us that *paradigmatic indispensability is not the same as permanence*. I claim only that, here and now, they are indispensable. It is, in fact, all any scientific concept can ever claim. The T-cell may one day be superseded by some deeper ontology we cannot yet imagine, perhaps a computational or informational framework that renders "cell types" as crude as "humors" now seem, or a systems-level understanding that dissolves discrete immune actors into dynamic network states. The T-test may give way to a radically different inferential framework - something as unimaginable to us as p-values would have been to Laplace. If so, we will not have been wrong to rely on them. We will simply have been doing the best science our paradigm allowed.

Ultimately, the point of this essay is not to resolve the realism debate but to unsettle a false certainty—to advocate, instead, for *epistemic openness*: the assumption that biological entities are obviously real while statistical tools are obviously not. If I have shown that both are abstractions, both are useful, and both deserve a kind of epistemic respect, then I have accomplished what I set out to do.

## **Conclusion**

What follows from all of this? If both T-cells and T-tests are load-bearing abstractions within our current paradigm, then the task of fostering mutual understanding falls on both sides of the divide.

For *biomedical training*, this means making explicit what is often left implicit: that biological entities like T-cells, receptors, and pathways are not brute facts handed directly to our senses, but operational definitions—constructs validated by their usefulness in predicting and intervening. A flow cytometry plot is not a photograph; it is an inference. Teaching this does not undermine these concepts; it deepens our grasp of their genuine power.

For *biomathematical training*—for statisticians, data scientists, and applied mathematicians—the lesson is symmetric. Mathematical concepts like the T-test, the p-value, or the Bayesian posterior are not mere games with numbers. They are tools for seeing patterns that would otherwise remain invisible. But they, too, are paradigm-dependent. The T-test made sense in a world of small samples and agricultural experiments; its role in modern precision medicine is real but not eternal. Teaching statisticians about the history and philosophy of science—about paradigms, constructs, and the contingency of even our most successful tools—would cultivate the same epistemic humility we ask of our biological colleagues.

What unites these two calls is a single insight: *paradigmatic indispensability implies epistemic openness*. If all scientific concepts are load-bearing abstractions within a paradigm, and if the history of science teaches that paradigms shift, then dogmatism about which abstractions are "really real" is intellectually indefensible. The philosophical position we have arrived at leads naturally to a practical stance: we must hold our concepts, whether biological or statistical, with sufficient humility to recognize their contingency, while taking them seriously enough to build careers and save lives upon them. This is not skepticism; it is the mature recognition that today's indispensables may be tomorrow's curiosities, and that the best we can do is work faithfully within our paradigm while remaining alert to the shifts that may come.

A journal club where the oncologist and the statistician see each other as engaged in the same essential act: using abstraction to extend the reach of human perception. A collaboration where the question "What is the mechanism?" and the question "What is the evidence?" are recognized as complementary, not competing. And perhaps, when mechanistic intuition and statistical evidence point in different directions, a pause—a moment of epistemic openness—before reflexively privileging the abstraction that merely feels more concrete.

The question "Is a T-cell more real than a T-test?" may have no answer. But the question "Are both indispensable to how we currently understand and treat disease?" has a clear one. Yes. And that is the question that really matters.

## References:

1. Fine, Arthur. 1984. "The Natural Ontological Attitude." *The Journal of Philosophy* 81(5): 261–277.
2. Hacking, Ian. 1983. *Representing and Intervening: Introductory Topics in the Philosophy of Natural Science*. Cambridge: Cambridge University Press.
3. James, William. 1922. *Pragmatism: A New Name for Some Old Ways of Thinking*. New York: Longmans, Green and Co.
4. Kuhn, Thomas S. 1962. *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
5. Laudan, Larry. 1981. A confutation of convergent realism. *Philosophy of Science* 48(1): 19–49.

6. Murphy, Kenneth, and Casey Weaver. 2016. *Janeway's Immunobiology*. 9th ed. New York: Garland Science.
7. Student. 1908. The probable error of a mean. *Biometrika* 6(1): 1–25.
8. van Fraassen, Bas C. 1976. "To Save the Phenomena." *The Journal of Philosophy* 73(18): 623–632.
9. Wigner, Eugene P. 1960. The unreasonable effectiveness of mathematics in the natural sciences. *Communications on Pure and Applied Mathematics* 13(1): 1–14.
10. Worrall, John. 1989. "Structural Realism: The Best of Both Worlds?" *Dialectica* 43(1–2): 99–124.