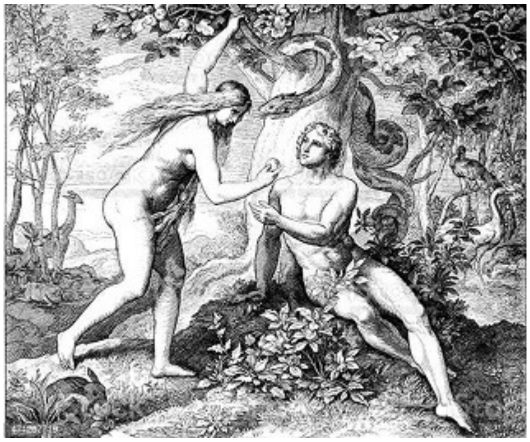

AN ATTACK ON THE SPIRIT OF KNOWLEDGE
HETERODOX RESEARCH SOCIETY

ALEXANDER VAWTER

While billions of tokens are generated and data centers heat up under sustained load, a parallel intensity and residual heat build in the minds of those engaged in popular discourse, where confidence about AI's role in discovery increasingly passes for inevitability. The future is not so much debated as it is collectively proclaimed, with scientific progress framed as something that will be automated, scaled, and rendered autonomous, finally freed from the slow, inconsistent, and all too human constraints of judgment and intuition. Sam Altman has stated plainly, "You will be thrilled that the AI has invented all of the science for you and cured disease and... made fusion work and just impossible triumphs we can't imagine." Such claims are delivered with a calm certainty that leaves little room for dissent, as if the nature of knowledge itself were already settled, awaiting only sufficient compute and execution. Increasingly, these claims are guarded by a moral sensitivity, wherein to suggest that humans ought to remain necessary, critical, or even simply present in the scientific process is framed as selfish, as though one were willfully constraining future medical breakthroughs or technological progress ("But the children!"). The resulting vision of a fused AI and scientific enterprise is one in which human involvement is minimized, and in our wet labs there could be robots who do the pipetting and the Bunsen burning instead of us flesh-bearing creatures. While the continued adoption and ubiquity of AI tools may indeed be inevitable, I do not believe the role these systems will play is equally predetermined. My contention is that the future of the laboratory will assume a more nuanced character, one in which humans and artificial systems stand in complementary relation, and in which the structure of discovery is such that humans remain uniquely suited to participate in it.

The Different Ways in Which We Contend with Nature

First, it is necessary to confront a simple but often neglected reality: not all scientific discoveries come about in the same way, nor do they require the same sorts of tools, reasoning, or personnel. The development of Einstein's theory of general relativity is of an entirely different character than the telescopic observations of distant supernovae that led to the discovery of dark energy. Even the Eddington expedition that confirmed Einstein's predictions differed profoundly in nature from the famous Einstein thought experiments themselves. Both were critically important and valid, yet some forms of discovery are often taken more seriously than others. Similarly, the unforeseen emergence of penicillin on Alexander Fleming's petri dish bears little resemblance to the long, deliberate pharmaceutical research process that led to the formulation of Percocet. These differences of experiment and discovery are not merely matters of scale or importance, but of structure. For this reason, it is useful to distinguish between two broad categories of discovery, which I will call ***incremental*** and ***transformative***. To be clear, these terms are not meant to imply magnitude or value. An incremental discovery may be vast in consequence, and a transformative one may be narrow or obscure. Rather, the distinction concerns how a discovery comes into being, the intellectual posture it demands, and the manner in which it reshapes our understanding of the world.

Incremental discovery tends to arise within an already accepted conceptual framework. It refines, extends, and optimizes what is already believed to be true. New measurements improve old models, new compounds are derived from known chemical families, and new effects are detected where theory already predicts their possible existence. This form of discovery is cumulative by nature and often collective in execution, relying on large teams, standardized methods, and increasingly sophisticated instrumentation.

Transformative discovery often begins with tension, anomaly, or dissatisfaction with the prevailing picture. Theories like general relativity did not emerge from faster calculation or denser data, but from a willingness to reconsider the underlying structure of space, time, and causation themselves. As such, it is a mistake to assume that all scientific activity proceeds in the same manner. Despite this reality, we have constructed a narrow vision of how one accumulates knowledge through “The Scientific Method” and related dogmas that bear an uncomfortable resemblance to the scholastic rules of the Middle Ages.

To this point, some of you may be thinking that I have simply redefined the difference between theory and experiment, and AI’s role in supporting the latter. To properly extinguish this assumption, I would note that even in the most practical settings, in the pure world of wet lab and experimental bench science, there exist many examples of transformative, intuition-driven discovery that cannot be captured by the logical or the incremental alone. Over many years, Barbara McClintock developed an unusually close familiarity with her work in maize genetics, following individual plants across generations while carefully drawing chromosomal changes by hand. In doing so, she began to notice recurring patterns that did not align with the prevailing view of genes as fixed and stable units. Certain genetic elements appeared to change position, interfere with other genes, and respond to stress in ways that made standard explanations feel inadequate. These observations would have been easy to dismiss as experimental error or contamination, particularly since there was no existing framework that could comfortably accommodate them. McClintock chose instead to continue working within this uncertainty, rather than forcing her observations to fit accepted models. The result was the identification of transposable elements, a finding that required geneticists to abandon the idea of the genome as a static structure and to treat regulation, mobility, and context as central features of heredity.

By contrast, the discovery of the Higgs boson at CERN followed a very different path, one that exemplifies a more incremental, utilitarian science. For decades, particle physicists had strong reasons to believe a particle such as the Higgs boson must exist if the Standard Model were to remain internally consistent. The task, then, was to detect an exceedingly faint signal buried within an enormous volume of particle collision data. Thousands of researchers collaborated across institutions to design detectors capable of tracking particle interactions at unprecedented energies, to calibrate those instruments with extreme precision, and to develop statistical methods robust enough to distinguish a genuine signal from background noise. The result emerged slowly through years of data collection and analysis, until the accumulation of evidence crossed the agreed-upon five-sigma statistical threshold required to justify a claim of discovery. This was science carried out through coordination, standardization, and scale, dependent on disciplined execution within a framework whose contours were already well understood.

Left Brain and Right Brain

If the distinction between incremental and transformative discovery feels unsatisfying, it can perhaps be reframed in terms more commonly used in such discourse. Much of modern science has become what I call ‘engineering-pilled’. It is oriented toward optimization and refinement within already accepted paradigms, well suited to scale, institutional coordination, and incentive structures that reward throughput, reproducibility, and predictable output. Once a field has largely exhausted the conceptual ascent available within a given framework, it tends to continue climbing along the same face of the mountain, extracting more precision and efficiency rather than questioning whether the route itself is worth following. In this sense, the dominance of engineering-pilled science today is a consequence of our institutional structure.

Running alongside this, though increasingly treated as suspect, is what might be called ‘philosophy-pilled’ science. This sort of inquiry is typically more concerned with orientation and interpretation than with optimization. It asks whether the assumptions that structure a field are justified, whether the questions being asked are the right ones, and whether the prevailing picture of the world is even adequate. In contemporary scientific culture, philosophy and metaphysics are often dismissed as distractions, tolerated at best as historical curiosities and scorned at worst as impediments to progress. This discomfort would have puzzled many of the figures now credited with the most profound reorganizations of our understanding of nature. Many rationalist sorts wish to claim that the religious or otherwise "misguided" efforts pursued by the vast majority of historically great examiners of nature were exactly that, misguided, and can be separated from their more serious works. However, this is to miss a feature for a bug inside the minds of great men.

James Clerk Maxwell's unification of electricity and magnetism emerged from a conviction that nature was intelligible and lawful, a conviction inseparable from his Christian faith. Werner Heisenberg, reflecting on the limits of scientific explanation (in a saying often attributed to him, though its exact origin is uncertain), remarked that at the bottom of the glass of natural science one does not find rationalist logic, Popperian falsification, or the scientific method, nor Neil deGrasse Tyson, but *God*. Whatever seemingly "loopy" commitments those who came before us may have had, it is undeniable that those scientists responsible for the deepest conceptual shifts were not embarrassed by questions of meaning, interpretation, or ultimate structure, and as such they did not cordon philosophy off from science.

When we exclude this philosophical dimension, the bones of the enterprise become brittle, and science atrophies as we starve half the body that gives it structure. Under such conditions, it is hardly surprising that artificial systems trained to optimize appear poised to replace the human entirely, because if science is reduced to refinement alone, then of course refinement machines will eventually outperform us. The more interesting question is what happens when science remembers that its history is shaped just as much by acts of reinterpretation as by acts of precision.

Knowledge, Itself

Now that a groundwork has been laid for understanding the diversity of scientific processes and discoveries, one might reasonably ask whether one category or the other, incremental or transformative, engineering-pilled or philosophy-pilled, is better suited to the human or to the artificial agent. Existing models, and those conceivable in the near term, can clearly support incremental probing within established datasets, theoretical models, and experimental landscapes. They can identify new proteins, classify stellar objects, surface errors in the literature, and accelerate the slow work of refinement with extraordinary efficiency. Even so, this activity remains in service of a human generator of questions, priorities, and interpretation. What remains far less clear is whether such systems can genuinely reinterpret the laws of physics, or even mathematics, upon which they themselves are built.

As previously alluded to, if knowledge is nothing more than verified facts, certified by trusted authorities, then perhaps we would indeed be approaching the end of the road for human scientists. But this is precisely the thin and impoverished picture of knowledge that dominates much contemporary thinking. It is a picture stripped of meaning, metaphysics, and orientation, reduced to laminated facts sold in bright orange and green textbooks for $199 by McGraw Hill. One can treat McGraw Hill as one's final source of truth, or one can open one's eyes to the vast and rich cosmos which is full of opportunity to make one's own determinations.

In a recent shared talk with my friend Riva Tez, she outlined the rich and layered etymological history of the word "knowledge," revealing that our modern usage reflects only a narrow slice of its older meanings. In ancient Greek, γιγνώσκειν referred not merely to factual awareness, but to a form of recognition that was spiritual and intuitive in character. In Hebrew, יָדַע carried an embodied meaning, denoting intimacy and lived experience, as in Genesis 4:1, "And Adam knew Eve his wife." The Latin scire, from which much of modern scientific language descends, narrowed this field toward facts, laws, and systematic truths. By the time of the Renaissance and Enlightenment, this trajectory culminated in the call to "Sapere aude," where courage became bound to what could be measured, proven, and systematized.

A useful reminder of how knowledge narrows rather than widens with time can be found in the way scientific traditions treat ideas they believe they have already defeated. In the seventeenth century, Robert Boyle's *The Sceptical Chymist* engaged alchemical traditions with care, reconstructing their strongest claims and criticizing caricatures of his opponents as intellectually lazy. By contrast, many modern chemistry textbooks reduce alchemy to a brief dismissal, offering little engagement with its internal logic or experimental sophistication. The issue is not whether alchemy was correct, but whether it was understood before being rejected. If we as humans increasingly fail to engage rival frameworks with seriousness and depth, it is difficult to imagine how the systems we train will learn to do otherwise.

To Reclaim,

If the plural nature of science is to reassert itself, it will do so not through automation alone, but through a broader cultural reframing of what scientific work is allowed to look like. As science decentralizes and methodological diversity becomes less institutionally constrained, the future of the laboratory may take on a more collaborative and augmented character. There could be literature-parsing assistants, automated documentation systems, and intelligent experimental aides that form something like a benevolent “Big Sister” in the lab, ultimately widening the researcher’s field of vision rather than replacing it. Such tools can enable methodological flexibility rather than suppression and rigidity, thus allowing multiple ways of knowing to coexist within an information-rich environment that traditional scientific tooling has often failed to provide.

Our tools have always changed how we see and what we can do. Tools have never rendered us obsolete on their own. Whether AI augments or erodes the spirit of knowledge will depend less on its capabilities than on the assumptions we embed in its design. If science is remembered not as a pipeline of facts, but as a rich, human way of contending with nature, then the future of the laboratory remains open, unfinished, and very much driven by *us*.

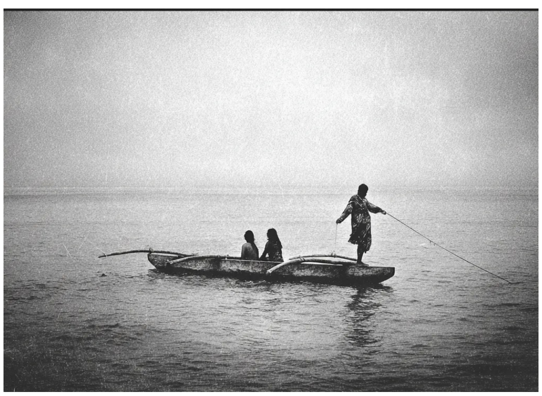

There is a vast, seemingly infinite beach,
Beneath its sands lie toys, treasures, and trinkets.
Some dig with shovels
Some have metal detectors
Some use their hands
But of these many,
There are few who look beyond the shores,
Spotting patterns in the tides,
Mapping the skies,
Extracting beauty from the mere perch of their gaze
Some discoveries can surely be made by force,
Yet others will only be made
If you are willing to look.