



The Ignorance Dividend

Comprehension is not the source of competence but its late-arriving, fragile byproduct—knowledge systems, from academic journals to AI models to financial instruments, have always certified outputs they don't fully understand. The current regulatory demand for transparency is a form of epistemic theater, and the alternative is governance built on enforceable liability, outcome monitoring, and consequence-tracing rather than on the fiction that complex systems can be made legible to those they affect.

In the spring of 1996, the cultural studies journal *Social Text* published a 39-page article by Alan Sokal, a physicist at New York University, titled “Transgressing the Boundaries: Toward a Transformative Hermeneutics of Quantum Gravity.” The article argued that physical reality is a social and linguistic construct, that quantum gravity carries progressive political implications, and that Rupert Sheldrake’s morphogenetic field constitutes a cutting-edge theory of structure. The editors published it without sending it to a physicist. Three weeks later, in the magazine *Lingua Franca*, Sokal revealed that every substantive claim was fabricated. The citations were real. The argument connecting them was nonsense—syntactically fluent, citation-dense, institutionally positioned, and propositionally empty. Read today, “Transgressing the Boundaries” is recognizable as something we now have a name for: AI slop. Text assembled from legitimate sources, stitched together with fabricated connective tissue, generating surface-level coherence without semantic content. Sokal, working by hand in 1996, produced the first piece of it. The editors certified it using the only tools available to them: the author’s credentials, the density of his footnotes, and the alignment of his conclusions with the journal’s intellectual commitments.

The affair was taken as a verdict on postmodernism. Nevertheless, that reading is 30 years old and it has exhausted its usefulness. What Sokal demonstrated with his crude audit—a one-time test for comprehension imposed on a system that had never required it—was something far more durable than the bankruptcy of cultural theory. It demonstrated the mechanism by which knowledge systems routinely certify outputs they do not comprehend. And that mechanism did not originate in the humanities but in earlier



systems of technical and institutional certification—finance, engineering, and the natural sciences—where outputs were routinely trusted before their mechanisms were fully understood. It is now the dominant operating logic of every computational system whose outputs matter.

When DeepMind’s AlphaFold—a deep-learning system designed to predict protein structures—predicted the three-dimensional structures of nearly every known protein, the researchers could describe the architecture but could not explain why the network arrived at a particular fold. John Jumper, who led the team, acknowledged as much: the knowledge was in the network, not in the team. When *Nature* published the results, the review process evaluated accuracy against experimentally determined structures. It did not require an explanation of why the network converged on its solutions. While *Social Text* certified on the basis of rhetoric and ideological alignment—prestige-tracking with no external check—*Nature* certified on the basis of predictive accuracy—truth-tracking by results. The difference matters. But what the two cases share is the absence of comprehension at the point of certification. *Nature*’s reviewers could verify that predictions matched known structures. They could not explain why the network produced them. The upgrade from *Social Text* to *Nature* is not a restoration of understanding. It is a more sophisticated proxy—better at catching fabrication, no better at producing explanation.

This is the argument that the transparency movement in artificial intelligence has not yet confronted. The demand for explainability assumes that governance requires comprehension of mechanism—that before a system can be legitimately deployed, someone must be able to narrate how it reaches its outputs. But what if knowledge systems have often operated without this narration for decades, and the absence of comprehension did not prevent them from functioning? It prevented them from being audited in one specific way. The question is whether that way is the only one available, or whether governance can be built on other foundations—traceability of consequence, contestability of outcome, enforceable liability for failure—that do not require anyone to understand the mechanism in order to hold its operators accountable for its effects.

The European Union’s General Data Protection Regulation (GDPR) enshrined a right to explanation for automated decisions. The AI Act layered additional transparency obligations on high-risk systems. The institutional consensus is overwhelming: opacity is the problem, transparency is the solution. But parts of the humanities operated for years under conditions of partial comprehension—circulating early work in postcolonial theory, queer theory, and science studies, fields whose central propositions were not fully legible to the institutional structures that would eventually absorb them. The editorial practice was not a failure of rigor. It was an infrastructure that allowed ideas



to circulate before they could be verified, because verification would have required the very disciplinary stability the ideas were challenging. The system's tolerance for incomprehension was the engine of its productivity. Sokal proved it could be exploited. He did not prove it should not exist.



The work produced under conditions of partial comprehension exceeds what comprehension-first approaches can reach. AlphaFold solved protein folding by learning regularities that 50 years of physics-based modeling had failed to articulate. Halicin was identified as an antibiotic because the network was unconstrained by the pharmacological categories a human researcher would have imposed. Daniel Dennett, the American philosopher of mind and cognitive science, spent decades building the philosophical framework for this observation. In *From Bacteria to Bach and Back* (2017), he argued that competence without comprehension is not an anomaly but the default condition of every productive system in nature—from evolutionary processes that accumulate design without a designer to computational systems that produce competence without insight. Darwin's strange inversion was the discovery that you do not need to understand how to make a perfect machine in order to make one. Comprehension is not the source of competence but its late-arriving, fragile byproduct—downstream, not upstream. The demand that AI systems explain



themselves before they are permitted to operate could be framed, in Dennett's terms, as a pre-Darwinian reflex: the insistence that competence requires comprehension, that the order of understanding must precede the order of production. Nature settled that question billions of years ago. This has a regulatory implication the current debate has not absorbed: if each generation of models diverges further from human cognitive categories—along the vector that Nick Bostrom's *Superintelligence* (2014) formalizes to its endpoint—then any regulatory architecture pegged to transparency weakens as its object becomes more powerful. Regulation must be independent of comprehension, built on enforceable liability and outcome monitoring rather than on explanations that are already unreliable and may soon be impossible.

The cost of this argument should not be concealed. In 2008, the global financial system detonated—and the standard reading is that the instruments were too complex, the derivatives too opaque, the risk models too arcane. Nassim Nicholas Taleb has argued in his book *Antifragile* (2012) that the problem with collateralized debt obligations and their synthetic derivatives was not that they were too complicated to understand. It was that many of them were built on risk models that assumed Gaussian distributions—bell-curve models that underestimate extreme events—in domains governed by fat tails, where rare but catastrophic outcomes occur far more often than standard models predict, or using formulas whose creators acknowledged they did not apply to the conditions in which they were being used. The instruments were not beyond comprehension. They were beneath it. But this is not where the damage originated. Finance had operated beyond the comprehension of most of its participants for decades before 2008—the Black-Scholes model, the options pricing revolution of the 1970s, the entire architecture of quantitative trading required expertise that most market participants, regulators, and certainly the public did not possess. That opacity was structural and, for long stretches, productive. What made 2008 catastrophic was not the opacity itself but the conditions surrounding it: the instruments operated in a regulatory vacuum, with credit default swaps traded over the counter and unregulated; the information asymmetry between the institutions packaging the securities and the pension funds, municipalities, school districts, and regional public authorities buying them was not incidental but engineered; the rating agencies whose certification the market relied on were paid by the institutions they rated; the compensation structures rewarded short-term risk-taking while externalizing long-term consequences onto populations who had no seat at any table where the risks were discussed. The prospectuses ran to hundreds of pages. The models were technically available. Opacity was destructive not in itself but because it operated without governance—without liability structures, without consequence-tracing, without enforceable obligations that bound the producers of incomprehension to




its outcomes. More legibility would not have helped, because the institutions receiving it had no capacity, no obligation, and no incentive to act on it. What was missing was not the ability to read the system but the obligation to answer for its consequences.

Sokal thought he was exposing a scandal in the humanities. In fact, he was drawing the first map of the territory we all now inhabit—a territory in which the most consequential systems produce knowledge that no one at any node in the chain fully comprehends, and in which the institutions responsible for certifying, deploying, and governing those systems have not yet developed a single practice adequate to that condition. Thirty years later, the institutional response to the permanent condition of operating beyond comprehension remains a demand for legibility—generating precisely the kind of epistemic theater Sokal set out to discredit. The ignorance dividend is real. It has always been real. It built the fields that now govern how proteins fold, how markets price risk, how ideas cross the boundaries that disciplines enforce. The question has never been whether to operate inside it. Algorithmic governance is a condition already being installed—in credit scoring, in predictive policing, in medical triage, in the allocation of public resources—and the institutions responsible for it have not yet built a single governance framework that does not depend on the fiction that these systems can be made to explain themselves. We need to prepare for a world governed by systems that will never be fully legible to the people they govern. Transparency is not the answer. It never was. The answer is governance that remains effective under enduring limits to explanation—or no governance at all.



BIO

Ilan Manouach is a comics artist and postdoctoral scholar in Computational Creativity at FNRS, University of Liège, with affiliations at Harvard metaLAB and Paris VIII. His work spans conceptual comics, tactile publishing, and AI governance, examining how computational systems reshape the conditions under which knowledge is produced, certified, and contested. He is the founder of [Echo Chamber](#), a Belgian comics research nonprofit, and [Topovoros Books](#), a Greek publishing house dedicated to translating key humanities texts. He is currently a Visiting Researcher at the Berkman Klein Center for Internet & Society, Harvard Law School.



Wild Papers editor: Ingo Niemann
Copy-editing: Rosanna McLaughlin
Illustration: Eva Fàbregas

Graphic design: Ana Domínguez Studio

© 2026, Ilan Manouach, Eva Fàbregas & Wild Publishing, a division
of Institute Art Gender Nature at HGK Basel FHNW, Switzerland