1. Good, Old-Fashioned Science

Dirac's 1931 prediction of the positron, and its subsequent experimental confirmation by Carl David Anderson, is a paradigm case of 'Good, Old-Fashioned Science' (GOFSci[1]). In 1928, Dirac had first published his eponymous equation, which sought to give a treatment of the theory of the electron consistent with both quantum mechanics and Einstein's special theory of relativity. Dirac thus was aiming at a unification of hitherto distinct phenomena, and in doing so, found that the resulting theory told him something entirely new, and entirely unexpected, about the world: that the familiar domain of particles is shadowed by a mirror-realm of *anti*-particles.

There is considerable power in this image. Guided solely by theoretical considerations, a singular mind was able to use daunting math to sneak a peak behind the veil, at the naked world beyond, to catch a glimpse of reality. Although Dirac himself probably never said so, the resulting theory has been claimed to contain 'most of physics and all of chemistry'. It is then little wonder that Dirac's remarkable discovery was to become the model for much of the progress in physics (and natural science *per se*) of the 20^{th} century.

And surely enough, the method of GOFSci continued producing impressive results: Pauli's prediction of the neutrino, Gell-Mann's and Zweig's prediction of quarks, the theory of the electroweak interaction, and finally, the Standard Model itself, completed in 2012 with the discovery of the Higgs boson.

But in more modern times, GOFSci seems to have run out of steam. The recent literature is replete with failed attempts at Grand Unified Theories (GUTs) producing unobserved proton decays; with axions, WIMPs, and superparticles that have failed to turn up in any search so far, despite the ever-increasing amounts of energy expended towards their detection; and with dark matter that continues resisting illumination[2]. All these were introduced on reasons that, at first blush, seem to parallel those of earlier, successful predictions. Yet, the detectors have remained silent. What gives?

In the following, I will propose that the method of GOFSci is reasonable only within a limited domain, and that extending it beyond that domain leads to the emergence of spurious problems, which then invite spurious solutions. Going forward, we may find ourselves forced to cast a critical light onto assumptions that have so far stood silently in science's shadow.

[1] In analogy to 'Good, Old-Fashioned AI' (GOFAI), the approach to machine intelligence mainly concerned with rule-based 'expert systems'.

[2] For an entertaining take on these issues, see the video in Ref. [1].

2. Three Dogmas of GOFSci (And Why We Should Renounce Them)

Dirac sought the theory of the electron in the unification of special relativity and quantum theory. As such, he was guided by a commitment to a thesis of *atomism*: that is, disparate physical phenomena of a given level find their explanation in a unified picture involving some more fundamental entity (an 'atom' in a generalized sense) at a lower level, where a simpler, more 'beautiful' theory might serve as a deeper explanation to open puzzles.

Guided by this atomistic principle, Dirac arrived at a theory that told us something new and true about reality: for consistency, there must be an entirely new type of 'anti'-matter: and lo, so there was. Thus, theory has the power of granting us insight into reality—science, at large, is about an objective, exterior world, and capable of delivering to us its naked truths. Call this second principle, simply, *realism*.

Finally, to achieve proper objectivity, the truths of science must be possessed of a certain absoluteness: it would surely not do for a scientific fact to vary from day to day, or person to person. To ensure this, the scientist must be fully separate from their discoveries, and abstract away from their own, limited perspective to have a chance at grasping truth unvarnished. The world can be appraised, in its entirety, within a single, universal perspective, a detached bird's eye *view from nowhere*—and only what is described in these terms has grounds to consider itself science proper.

GOFSci, then, consists of *atomistic* theories, simplifying phenomena to ever more fundamental levels, understood in a *realist* manner as directly speaking to the truths of the world, dispassionately assessed from an infinitely removed vantage point, a *view from nowhere*.

Of course, put so starkly, the dogmas of GOFSci amount to a bit of caricature—few would endorse them in their most naive form, certainly not explicitly. But in order to see our target clearly, it is helpful to paint it in bold colors, the better to pick it out from the surrounding weeds.

2.1. Realism and Competence without Comprehension

Scientific realism is the attitude that our best theories and models *get something right* about the world—what they tell us is not merely true, but true of an objectively existing external reality. Thus, it is rational to believe in the entities postulated by successful, mature scientific theories, even if those entities—like atoms or electrons—are not directly observable.

This is a reasonable attitude: after all, such belief is often rewarded—see the eventual discovery of the positron as predicted. Thus, the simplest argument for scientific realism is: it's true because it works. More elaborately, our best current theories must get *something* right, otherwise, our success in applying them (for prediction, or creation of new technologies) would lack any justification, and hence, be effectively miraculous (Putnam's 'no miracles'-argument [2]). Theories must faithfully represent something 'out there', otherwise, the competences they afford us seem inexplicable. Think about finding your way through a cluttered room in the dark: unless you know, at least approximately, where things actually are, getting through without knocking something over would seem an extraordinary feat of luck.

In other words, our *competence* in navigating the world seems to presuppose *comprehension* of its workings. But things are not quite so straightforward. Even *false* representations can grant outward competence. An example is provided by science-fiction author Douglas Adams: in a speech titled "Is There an Artificial God?", he discusses the practice of *feng shui* [3]. According to Adams, feng shui entreats us to arrange a living space in such a way that a dragon could move through it with ease. This, he rightly points out, is nonsense: there are, after all, no dragons. But nevertheless, it's entirely possible that, because we're skilled in thinking about how organic beings work, envisioning how a dragon would move through a living space results in something that's actually pleasant to inhabit. But if so, that doesn't give us any cause to believe in the existence of dragons.

The point can be made more formally. Cognitive psychologist Donald D. Hoffman proposes that 'fitness beats truth': an organism capable of perceiving 'reality unvarnished' will, under general conditions, lose the evolutionary race against one capable of distinguishing fitness affordances in the environment. Within the setting of evolutionary game theory, this leads to a rigorously provable theorem [4].

As an illustration, suppose that there is one organism capable of perceiving the absolute amount of some substance (perhaps a nutrient). Suppose that too little, as well as too much, of that substance is fatal. Then, that organism will be at a disadvantage compared to one that just perceives whether the substance is in the 'Goldilocks' zone of *just the right amount*; but the latter organism, unlike the first, does not represent a straightforward truth of the environment.

Hoffman contends that this implies that our perceptions—and by extension, experimental data and the theories built on their strength—ought to be regarded as *interfaces* representing ways to interact with the world, rather than representations of the world itself. While providing us with useful means to navigate our daily lives, we take their representations literally at our peril.

Developments of recent months have given us grounds to further doubt our direct line to truth. Although the true impact (if any) of large language models (LLMs) like ChatGPT on philosophy, like on virtually any subject, is yet to be determined, one important lesson seems to be just how far what the philosopher Daniel Dennett [5] calls 'competence without comprehension' can take you.

There is no need to delve too deeply into ChatGPTs capabilities here—that ground has been well-trodden in the short time since it was first announced. Suffice it to say that it essentially works like the predictive text feature on modern cellphones—albeit in a much more sophisticated manner:

given a string of text, it seeks to predict the best possible continuation. It does this, as anybody who has spent some time interrogating it, shockingly well, to the point that its responses are often hard to distinguish from those of a human interlocutor.

However, in this prediction task, what a given word *means* simply doesn't enter[3]: it has learned statistical regularities of large quantities of text, and uses these to assemble abstract tokens. In other words, it doesn't merely have a *misleading* or *incomplete* representation of the outside world, it has *no representation* of it at all—all it represents (if it represents anything) are statistical relations over words. In short, it is just the sort of miracle Putnam could not conceive of.

2.2. Farewell to the Atomistic Age

The dominant strategy of (particle, at least) physics, when it comes to finding explanations of phenomena, has been: let's crack it open to see what's inside. Not literally, of course: there is no real sense in which particles produced in a collision were 'contained' within the colliding entities. But the idea here is that what's 'inside' may have explanatory power: that what's puzzling on the larger scale is explained by looking deeper. The discovery of chemical elements and their ordering within the periodic system reduced a plethora of substances to arrangements of a few dozen atoms; protons, neutrons and electrons then further curtailed the remaining complexity.

Accordingly, we may term this approach *atomism*: the idea that apparently arbitrary and puzzling features at one level of description can be reduced to entities at a lower level that, importantly, are generally considered *simpler* in some appropriate sense. It's the latter proviso that delineates this notion from *reductionism*: the second law of thermodynamics can be *reduced to* the statistics of molecules, but their behavior is, in general, of vast complexity, which gets washed out in the average. By contrast, the reduction of the second law to properties of some simple substance, like phlogiston, would have been atomistic in this sense.

Moreover, the more fundamental theory ought likewise be more general, able to extend to novel phenomena not part of the original layer. Dirac's theory of the electron is atomistic in this sense: it yields a fundamental explanation for the earlier, phenomenological Pauli-model, and extends it in two senses, to relativistic regimes, and to the entirely unsuspected realm of antimatter.

But the atomistic *Ansatz* rests on a misconception: that parts are simpler than the whole. While this is true of many everyday systems, and thus, has some intuitive heft, it turns out to be false in a precise sense. To see how, we need to be explicit about what is meant by the intuitive notion of *simpler*.

A suitable notion of complexity is provided by algorithmic information theory [6]. The *Kolmogorov complexity* of a given (computable) object, such as a string of symbols, is given by the

[3] For an argument to that effect see Appendix A.

shortest program that produces this object[4]. Thus, a simple string, like '010101 . . .', produced by a short program like 'print '01' n times', has a low complexity, while an unpredictable, essentially random string will be highly incompressible.

The atomistic strategy is equivalent to trying to find a simpler program reproducing all known data (and ideally, making novel predictions). Success means explaining more with less. But this faces a dilemma: either, we eventually find some fundamentally simple program. Then, its remaining complexity will be without any possible further explanation—irreducible to anything else, a 'brute fact' of the world. Or, the process of reduction carries on, with ever deeper levels: in both cases, a final answer never comes.

Furthermore, how do we know when to stop? The Standard Model, our current best theory of elementary particles, contains something like 23 unexplained parameters. Are these just part of the input to the smallest program, or do they have a further explanation? And even if they do: what is won by just iterating to the next more fundamental level with *its* unexplained input parameters?

We will leave this issue for now, but not before surreptitiously hinting at a possible answer: the Kolmogorov complexity of the whole can be quite a bit *less* than that of any of its parts. To see why, we need only take a look into Argentinian writer Jorge Luis Borges' *Library of Babel* [7]. In it, every single book of 410 pages length is stored. In general, thus, *each single book* may contain a great amount of information (including, for instance, a complete account of the final rung of the atomistic ladder), but the *whole library* is completely—and more importantly, *concisely*—described by the preceding sentence.

Information is not the sort of thing that piles up if you collect it, like bricks of Lego do. The right pieces of information, if combined, may amount to little, if anything at all, no matter how vast their individual complexity.

2.3. From Nowhere to Now, Here

The scientific observer is a detached entity, capable of surveying the world from an objective vantage point, impartially cataloguing experimental facts. What's true for them is *true*, simpliciter. By postulating this *view from nowhere*, GOFSci seeks to ensure that the scientist can be removed from the business of science, permitting a focus on objective reality freed from any subjective tilt. But the world is more complicated than that—it is too rich to fit within any single perspective.

Consider the raw clay of science, the experimental data: usually, data are thought to represent impartial facts of the world against which theories are tested, to be upheld or tossed upon the scrapheap of history. If your theory says the detector should click, and the detector doesn't click,

[4] Since any (universal) computer can simulate any other, this notion won't depend on the computer running the program, except perhaps for some constant overhead.

then nature has rendered her verdict.

This naive picture runs into well-known issues. As Einstein pointed out to a dismayed Heisenberg, "it is the theory which decides what we can observe" [8]. After all, how should we know what the detector's click indicates without a theory of how the detector works?

But the problems run deeper. Ultimately, all data are subject-object correlations: the subjective experience of an audible 'click' of the detector is taken to indicate a certain event—say, the decay of a particle, 'out there' in the world. But a correlation on its own does not convey any information. Suppose I have two colored sheets of paper, a red and a blue one; and I slip each into an envelope. There is a perfect correlation between the slips, but that alone tells me nothing about the color of either. Only if I open up one envelope and find the red paper can I say that the other one must contain the blue slip.

Drawing information from a correlation requires 'fixing' one of its constituents. To be able to think of data as objective 'facts of the world', the subjective 'pole' has to be taken for granted—to be universal for every possible observer. But what grounds could we have to do so? After all, what goes on in *your* head is not directly available in *my* world.

But the greatest challenge for the one-size-fits-all conception of the world in GOFSci comes from quantum mechanics. For nearly a century now, it has been widely felt that the quantum revolution entails a profound departure from classical concepts. However, there has been little consensus on where, exactly the two part ways: Schrödinger saw the essential novelty in *entanglement*; Feynman pointed to *interference*. But in recent times, a certain kind of *perspectivalism* has emerged as a plausible candidate.

The world of GOFSci is one where, in principle, everything can be brought under a single umbrella, told in a single story. A precise formulation of this intuition goes back to the father of modern logic, George Boole, who in 1862 formulated what he termed the 'conditions of possible experience': 'When satisfied', he proposed, 'they indicate that the data *may* have, when not satisfied, they indicate that the data *cannot* have, resulted from actual observation' [9].

These conditions come in the form of constraints that the probabilities of observed events must fulfill. Today, Boole's conditions (a certain subset of them, more precisely) are more famous under the name *Bell inequalities*[5], and their violation—indicating the reality of what Boole thought a conceptual impossibility—was honored with the 2022 Nobel prize in physics.

Why was Boole so certain that observational data could never violate his conditions? For their derivation, it suffices to assume that (1) all observable quantities have definite, unique values which are (2) independent of the act of observation. In other words, Boole's conditions codify the ideal of the detached scientific observer. If they are satisfied, it is possible to reason about the world from

[5] A connection first pointed out by Pitowski [10].

within a single logical system (a *Boolean algebra*). Their violation in quantum mechanics thus implies that the world can only be apprehended by means of a collection of perspectives 'glued together', but is not capable of being subsumed within a single, unified frame: a congeries of views from 'now, here' rather than a single view from nowhere.

3. Interstitial Realism: The World as Interface

Should the kind reader have followed my exposition up to this point, they may be left somewhat troubled: if science does not serve to uncover ever more fundamental, simple truths about reality, independent of any individual perspective—then *what*, in the end, *does* it do? Should we chuck out the whole enterprise? Accept a fundamental relativism, where anything goes and your truth is as good as anyone's?

Luckily, there is a way for science to keep what works, and throw out accumulated excess baggage (all the while revealing some thorny old problems to be mere artifacts of unquestioned assumptions). We've already encountered a few hints to guide us: if science no longer allows us to view nature uncovered, it can still aid us in finding the best way to engage, to *interface* with it; if we can't look down to ever more simple fundamental atoms, we can instead turn our gaze upwards, at the whole instead of the parts; if there is no single, all-encompassing story to be told, we can instead learn to appreciate the plenitude and diversity of this glittering jewel of perspectives.

To develop these hints, let's start, as all good science should, with an experiment. Find a sheet of plain, white paper. Note that it has a highly concise description—and consequently, a very low algorithmic complexity. Then, place two hands firmly beside each other on one side of the paper, gripping it between thumb and forefinger; and start, slowly, to tear down, until you have two pieces in your hands. Now put one off to the side, and focus on the other: try to find a complete description, as you did when the paper was whole. You'll find that, all of a sudden, what you have created is much harder to describe. In other words, you have created information—if not quite, perhaps, *ex nihilo*.

This bears repeating: information can be created by *separation*. By dividing the paper into *this* part and *that* part, you have brought something new into the world that wasn't present before: you've made it a tiny bit more complex. Suppose we were to try and explain this complexity using the atomistic methods of GOFSci: we could find no simple account of the paper's torn edge. But now, take the other part, the one discarded: it fits the one in your hand perfectly—it contains exactly the same information, and its shape mirrors that of the other. Reunited, each explains the other fully, and the apparent complexity is sublimated: no irreducible 'brute fact' is left.

I contend that this is the *only* way information enters into the world. To the extent that we see a world of nonzero information content, this information was created, likewise, in a process of separation. But what has separated? Suppose you were a tiny ant, living on the sheet of paper. Once it is ripped, you suddenly find yourself inhabiting a world of high complexity—which you could then study with the tools of science. But what makes it appear thus is really only your limited perspective: if you could transcend it, observing the undivided whole, you could again describe it simply. Nevertheless, the information thus created is—for you—perfectly real.

The separation of the world into distinct perspectives is what creates the complexity we see. But we are not limited to being mere passive observers of the information created at the interstices of the different perspectives of the world. Recall that the failure of Boole's conditions entails the impossibility of the detached observer: thus, this information must depend on how we probe it—how we interact with it. This is a boundary to be pushed—and prodded, and fiddled with, and manipulated, to investigate how it reacts. It is, in the metaphor of Hoffman, an interface between you and the rest of the world, residing at the interstice of separated perspectives.

4. Pushing the Boundaries

The above proposes a wholesale revision of the basic metaphysical assumptions typically thought to form the basis of scientific investigation. As such, it should not be accepted without compelling reason.

In recent years, it has been widely perceived that science is, in one way or another, in crisis. Whether physics has become "Lost in Math" [11], or the fraction of disruptive research has been steadily declining [12], a new way forward seems to be called for. So what can the view championed here offer?

Many of the current problems in fundamental physics take the form of questions asking for the particular value of a parameter: particle masses, the relative strength of the various forces, the cosmological constant, to name a few. These ask, in the end, for the minimum information needed to fully describe the world; and atomistic GOFSci looks for it in ever more fundamental entities further down the ladder—thus, creating a long list of predicted, but so far unobserved particles and phenomena.

But the view proposed here invites us to take a different stance. The complexity we see in the world is a consequence of apprehending it from a limited perspective; hence, that there appears to be some irreducible information to the description need not bother us. It is just the way the paper tore, that jagged shape describing either side of the split.

Still, this seems to suffer from the same old problem: how do we know when to stop? But here, we can avail ourselves of some new tools: unlike GOFSci digging ever deeper in hopes to strike explanatory gold, a boundary can be approached from two sides. Thus, where we have so far only looked from the outside in, we can now look, too, from the inside out.

How much information does it take to specify the state of the observer? Describing the boundary from one side is, after all, equivalent to describing it from the other; thus, the information we see in the world must be lower-bounded by that needed to yield a mind capable of this observation. Then, complementary to probing the boundary from outside with ever simpler theories, we can approach it from within, as well, looking for convergence in the limit.

Interestingly, efforts to quantify the information in a mental state using radically different methods, by the cosmologist Max Tegmark [13] and the philosopher Tor Nørretranders [14], have produced similar values—a few tens of bits. If this is true, we might expect a 'final theory' to have a comparable complexity. This remaining information will not be a mere 'brute fact' about the world to be accepted, but rather, emerge dynamically from the separation of perspectives.

Furthermore, by giving consideration to both sides of the boundary, the individual's perspective, their lived experience, need no longer be something to be abstracted away and replaced with some imaginary standard, or universal—and in practice all too often exclusionary—point of view, but may itself be the object of scientific investigation. Here, the mystery of consciousness looms large.

. . . .

A final simplicity will not be found below, in the fragments of the world, but only in their union. We should no longer think the world as assembled from a great many tiny Lego pieces, but rather, as fragments of—better, perspectives on—a single, underlying whole [15]. More, here, is less—less complexity in want of explanation, that is. Ultimately, the information content of 'everything' tends to zero—and it becomes hard to tell apart from nothing [16], from which emerges *something* at the interstices of individual perspectives.

GOFsci was only ever an approximation, reasonable where we could neglect the distinctions between different perspectives. Brushing up against its boundaries, a new approach is called for. I have proposed that this may be found in an *interstitial realism*, where information emerges along the fault lines of distinct perspectives.

This points us in a different direction in the search for a more fundamental theory, away from the latest microscopic Lego pieces, and towards more all-encompassing notions. In this sense, the final theory may well be the ultimate ensemble theory [17], describing the world in the limit as the collection of perspectives on it.

Where does this leave us? If the above has any plausibility, it is perhaps rather through the limits identified, than through the sketch of the way forward provided. Every perspective is bounded, even that of modern science, won, as the quote has it, by standing on giant's shoulders. What lies beyond can only be speculated about—unless we leave our exalted vantage point, climb down, find our way through the weeds to the boundary's edge—and push.