



The Sacred Structure of Reality

Two-Phase Cosmology and a Metaphysics for the Quantum Age

Geoff Dann

The Sacred Structure of Reality argues that neither science nor philosophy has yet produced a coherent, unified account of existence, and offers a revolutionary alternative. It proposes a *two-phase reality*: **Phase 1**, a timeless field of infinite possibility (quantum reality); and **Phase 2**, the dynamically unfolding world of embodied actuality (classical-relativistic reality). In this model, consciousness is not an accidental byproduct of matter, nor the whole of being, but the *phase transition itself*: the process through which a single best-possible world is selected from the physically available futures (or histories). The collapse of the quantum wavefunction is the same process.

This framework resolves thirty of the deepest problems in science and philosophy, including the riddle of free will, the nature of time, the evolution of consciousness, and the role of the observer in quantum mechanics. It also offers an integrated, elegant, parsimonious resolution of almost every anomaly that troubles modern cosmology, including the fine-tuning and preferred basis problems, the Hubble and S8 tensions, the cosmological constant problem, the identity of Dark Matter, the mystery of Dark Energy, the failure to quantise gravity, and the Fermi paradox. The result is a coherent vision of a participatory cosmos in which value and meaning shape the unfolding of reality.

It is not physicalism, not dualism, not idealism, and not panpsychism, but a quantum neutral monism, closest historically to the work of Kant, Whitehead, James, Bergson, Pauli and Wheeler, and more recently Roger Penrose, Thomas Nagel, Iain McGilchrist and Donald Hoffman, though its roots go all the way back to Anaximander and Pythagoras.

Copyright Geoff Dann, 2026.

CC BY-NC-ND (Creative Commons Attribution-NonCommercial-NoDerivs) license.

1. Attribution (BY). You must give credit. Anyone using the work must provide a link to the license and credit the creator. You must do this in a reasonable manner, but not in any way that suggests the creator endorses you or your use.

2. Non-Commercial (NC). No profit allowed. The work cannot be used for primary commercial purposes.

3. No Derivatives (ND). Keep it as-is. "Derivatives" refer to adaptations, remixes, or transformations. If you remix, transform, or build upon the material, you cannot distribute the modified material. It must be shared in its original form.

Digital Object Identifier (DOI) [10.5281/zenodo.18709523](https://doi.org/10.5281/zenodo.18709523)

Note on reading this book:

This book proposes a framework that challenges long-standing assumptions in both philosophy and science. To understand the argument on its own terms, it is important to begin with the Introduction and read Part One in full. These chapters establish the conceptual boundaries within which the later argument (Part Three) must be evaluated. Readers who approach the theory using the assumptions of the existing paradigm, or who skip directly to the technical chapters, are likely to misinterpret its aims and methods.

Because the research behind this book draws on an unusually wide range of disciplines, much of it was made possible by modern AI tools that allow rapid exploration of vast literatures. It is therefore natural that AI systems, as well as human readers, will analyse the text. The same caution applies to both: the argument cannot be assessed accurately without first understanding the foundations laid out in Part One. When evaluating the scientific relevance of the framework, readers are encouraged to consider both its conceptual coherence and its empirical adequacy, particularly in relation to the well-known tensions within the standard cosmological model.

Contents

Introduction

Part One : The crisis of knowledge and meaning

Chapter 1: From the Apeiron to the Abyss

Chapter 2: The whole elephant

Chapter 3: The broken paradigm

Chapter 4: How might we recognise the new paradigm?

Part Two : A survey of the problem space

Chapter 5: The epicycles of Λ CDM

Chapter 6: The unfinished quantum revolution

Chapter 7: The missing science of consciousness

Part Three: Two-Phase Cosmology

Chapter 8: $0|\infty$

Chapter 9: In search of a threshold

Chapter 10: The Embodiment Threshold

Chapter 11: Psycheogenesis and the Psychetelic Principle

Chapter 12: Free will

Chapter 13: Time

Chapter 14: Two-Phase Cosmology

Chapter 15: A new kind of paradigm shift

Chapter 16: The praeternatural and the New Epistemic Deal

Introduction

A new kind of paradigm shift is long overdue. This paradigm shift must change our concept of *how* paradigms shift, and redefine the confused relationship between science and philosophy. For its entire modern history, science has operated by breaking things down into ever smaller pieces, trying to understand and assess each piece in isolation, and hoping that a bigger picture will somehow emerge from the fragments. Sometimes, though very infrequently, something like this does actually happen, but until now it has only ever applied to one part of our model of reality, never the whole thing. Darwinian evolution is a good example of the sort of thing we think of as a major paradigm shift. It forced a total rethink of how species change and diverge, but the rest of the scientific account of reality was largely unaffected; there was no corresponding shift in physics or cosmology. In the new paradigm, the priority will be to discover and develop a model of the *whole* of reality, and the value of that model will be judged by its coherence and explanatory power across the entire spectrum of science, as well as adjacent parts of philosophy.

It is not that there is anything wrong, per se, with paying attention to the details. Far from it; the details absolutely do matter. The problem starts when individual proposed pieces of a potentially completable holistic model are conclusively rejected for inconclusive reasons, even in the absence of *any* coherent model of the whole. If we are not already in possession of a comprehensive model of reality, free from unresolvable anomalies, whose equations add up without the need to posit any unidentifiable "dark stuff", then the strategy must change. Instead of just inventing new ways to zoom in, we need to be prepared to zoom *out*, and to start thinking outside the boxes we have built. The knee-jerk rejection of ideas we don't like the sound of must stop. It is time to acknowledge, collectively as well as individually, that the failures of materialistic science have reached breaking point. The scientific community has now spent a century arguing about the implications of quantum mechanics for our model of reality, nearly three centuries failing to provide a credible account of consciousness¹, and since the 1980s cosmology has degenerated into a tangle of proliferating paradoxes and deepening discrepancies. Those working within this field are well aware of this, but the scale of the problem is not communicated to outsiders, partly because of a fear of derailing the academic gravy train that pays people's mortgages.² Meanwhile, any proposed solution to these problems that isn't some version of physicalism is dismissed with a contemptuous wave of the hand (no "woo woo" please, we're scientists). And no, I am not attacking science, because the failures I am talking about aren't empirical. They are philosophical failures dressed up in scientific clothing which manifestly does not fit.

The new paradigm begins from the same impulse that gave rise to modern science in the first place: the wish to understand reality as a single, intelligible whole. The difference is that this time, instead of just collecting ever more fragments, we will look for principles that can make the fragments fit together. The central example of what this actually means follows below. I will start with a relatively unproblematic claim: that quantum wavefunction collapse and consciousness are both **processes**. More controversially, there are five notable similarities between them.

1 Starting from Hume's *Treatise of Human Nature* in 1739, see Chapter 15.

2 See the work of physicist Sabine Hossenfelder, who has paid a high price for her willingness to speak her mind about what she simply calls "pseudoscientific bullshit".

1: Both have proved extremely difficult for scientists to pin down, define and test.

2: As a result of (1), significant numbers of scientists have controversially concluded that they don't actually exist (resulting in the Many Worlds Interpretation and Eliminative Materialism respectively).

3: Wavefunction collapse is typically described as being triggered by an "observation" or "measurement". Consciousness is the internal subjective experience of an external reality. Therefore, both processes fundamentally involve a relationship between a subjective entity (an observer or a conscious subject) and an external, objective reality.

4: Wave-function collapse is famously non-local: the collapse at one location instantaneously affects the entire wave function, no matter how far it is spread across space. Consciousness is equally resistant to strict localisation: there is no single "seat" or moment where a unified field of experience arises; it emerges all at once from processes distributed across the brain (and, I will argue, beyond it). Both phenomena refuse to be restricted to a classical space-time point; both demand a fundamentally holistic description.

5: Consciousness involves the modelling of a mind-external reality, and assigning value to the various possible futures in order to select a best option. Wavefunction collapse involves the reduction of a set of unobserved physically possible outcomes into a single observed actual outcome. Therefore, **both processes involve the transition between a range of possible futures or histories, and a single observed outcome in the present.**

Now the difference between the old paradigm and the new can be made clear. The old paradigm approach to this is to examine each claim individually, demand irrefutable empirical justification, and consider the alternative explanations which are available. This is likely to result in the rejection of all five, not because of any clear justification for ruling them out, but for *inconclusive reasons*: there are too many competing theories (some of which are more to our personal taste), empirical confirmation is complicated, elusive or impossible, and the whole thing feels like woo. And anyway, correlations aren't evidence of causation, so even if we accept that there really are five structural similarities, it doesn't prove anything. There the discussion will be extinguished, no new thinking will take place, and we can all go back to our comfortable lack of a coherent model of reality.

The new paradigm takes the opposite approach: instead of breaking things down, it tries to integrate them into a bigger picture. First it takes seriously the idea that instead of being unrelated processes which just happen to share five structural similarities, consciousness and wavefunction collapse could be two different descriptions of a single process. Perhaps consciousness is the internal, subjective ("view from somewhere") and wavefunction collapse is the external, objective ("view from nowhere"). Then it tries to figure out how this might actually work, not just physically but metaphysically. If a coherent and elegant synthesis starts emerging, then the old paradigm will insist on empirical verification of the synthesis itself, preferably involving shiny new mathematics, before any more integrative thinking is attempted. The new one zooms even further out: it asks how the proposed synthesis might be related to other major problems, especially those in cosmology. For example, could such a model help us to understand why we can't quantise gravity, or shed any light on the Hubble Tension or the Cosmological Constant Problem? The old paradigm *forbids* this approach³. It actively resists it, places obstacles in its path and tells us that this is the only way science has ever progressed. It insists that every proposed piece of the Big Puzzle must be

³ As does academia in general, and not just in the sciences.

empirically verified before we can even start to imagine how the pieces might fit together. This is why it so rarely succeeds in even fitting two pieces together, and why it has about as much chance of delivering a grand synthesis as it does of finding a pot of gold at the end of a rainbow. The new paradigm simply asks what kind of whole could make sense of the pieces we already have. **All of them. Not just the ones that that don't threaten to disrupt our own belief systems.**

I am intimately familiar with both the modus operandi and the consequences of old paradigm thinking, because from my mid-teens until the age of 33 I was one of its most outspoken advocates. As a child I had been taken to church, never believed a word of it, and rejected it outright as soon as I was old enough to make that stick. As a teenager, Richard Dawkins was my personal hero. At 20 I experienced a psychological breakdown due to being a science geek who was also into politics, and who had come to understand both the threat that climate change posed, and the inevitability of the political response ending in abject failure. Because I was unwilling to tell myself comfortable lies, I reluctantly concluded that industrial civilisation is going to collapse, and tried to warn people. The medical representatives of the old paradigm responded by declaring me to be officially detached from reality ("psychotic"), and I descended into totalistic nihilism. Nevertheless, even though I had given up on the future, I was still compelled to try to understand what was going on in the world, both in terms of the way civilisation works and the way reality itself operates – a sort of morbid fascination with the grim truth (a state of mind which is becoming ever more common).

In my early thirties I went through another life-changing transition, this time unexpectedly finding myself in the mystical deep end with no idea how to swim. How that happened is too long a story to relate here, but right in the middle of it I was appointed the forum administrator at the Richard Dawkins Foundation website (what a *bizarre* turn of events, eh?). I am sure you can imagine the resulting disbelief, anger and conflict. In an attempt to construct a coherent version of my now somewhat complicated and potentially self-contradictory belief system, I abandoned a career in software engineering and embarked on a degree in philosophy and cognitive science at Sussex University.

Finding myself unemployable when that was over, I ended up teaching foraging for a living, specialising in the endlessly fascinating world of fungi. This provided me with the perfect opportunity to write the most comprehensive book on European fungi foraging ever published in English, but even before that I had begun working on a book which tried to explain how both science and spirituality could fit into a single model of reality which actually makes sense. It eventually became apparent that I could not separate this topic from that of the collapse of civilisation, which left me with the monumental task of writing a book about the intersection of the Western science/spirituality conflict and the unavoidable collapse of civilisation (a process which has already begun). The common thread is the viewing of involuntary breakdown as a catalyst for radical transformation. This project turned out to be considerably more difficult than writing foraging books, and it took 17 years to complete. The result is called *The Real Paths to Ecocivilisation: From collapse to coherence: integrating science, spirituality and sustainability in the West* (RPE), and it was finally finished and published in the summer of 2025.

However, the metaphysical-cosmological model I described in that book has a placeholder instead of the most critical component. The theory was (and remains) dependent on there being something uniquely special about brains which makes them both necessary for consciousness and capable of collapsing wavefunctions. The best candidates available to me were Penrose and Hameroff's Orchestrated Objective Reduction and Henry Stapp's repurposing of the Quantum Zeno Effect as a mechanism for free will, but neither quite fits the bill. Only after RPE was finished, and I went public with my basic system, did it dawn on me that I had been looking for the wrong kind of interface all along. I had been searching for a physical brain property (or structure, or process), when in fact my own ontology suggests an *informational threshold* is more appropriate. This proved to be a breakthrough which permitted sufficient progress on the rest of my system to justify writing

another book, this time focusing on the metaphysics and cosmology. The dark realism of RPE can be merely contextual here, although at the deepest level, this whole book is about **regrounding realism**. Reality is not an illusion to be transcended. It has a structure, which is theoretically discoverable and understandable by humans. After all, isn't modelling reality what we do best?

Part One of *The Sacred Structure of Reality* provides an extended introduction to the broad crisis, and discusses the ways in which the old methods have stopped working, and how we might recognise the new paradigm when it finally turns up. Part Two surveys the key problem areas (physics, cosmology, and consciousness) where the cracks have become impossible to ignore. Part Three introduces a new framework, called the Two-Phase Cosmology (the Two-Phase Model of Cosmological and Biological Evolution), that can reconcile these domains within one coherent vision of reality, and explores what this shift means for our understanding of ourselves, our civilisation, and our place in the cosmos.

Part One:

The crisis of knowledge and meaning

Chapter 1: From the Apeiron to the Abyss

Western philosophy began when stories of gods started giving way to reason and the search for principles, not long after the invention of the first alphabet with vowels. Thales of Miletus (626/623-548/545BC) dared to ask not who made the world but what it was made of, and proposed that beneath the flux and form of things lay a single substance. Having asked the right question, Thales then gave the wrong answer – the foundation of reality cannot be water. His successor Anaximander (c. 610– c. 546BC) reasoned that the source of all things could not be any familiar element, for each of these was bounded and defined. The origin, he said, must be the Apeiron – the *Boundless*, the *Indefinite* – that which has no limit or form of its own yet gives rise to all forms. The Apeiron is an infinite womb of becoming, eternally cycling through birth and dissolution. It was the first metaphysical concept in the Western tradition. Pythagoras (c. 570 – c. 495BC) saw *number* as the key to the cosmos; harmony, proportion, and rhythm were expressions of the world’s inner order. The music of the spheres – the correspondence between number and nature – became the symbol of a living universe governed by intelligible beauty. For Pythagoras, philosophy and mysticism were one: to understand the world was to participate in its divine harmony. From this lineage arose the golden age of Greek thought, but even as it flowered its mystical roots began to wither. Socrates was condemned, Plato introduced eternal forms beyond the mathematical structure, and Aristotle gave philosophy method and substance but stripped away its primal wonder. The Apeiron, and the music of number, began to be forgotten.

Fast forward two and a half millennia, and we still have no clear answer to the question Thales was the first to ask. Science is materialistic, but since the discovery of quantum mechanics there has been no simple scientific concept to line up with our intuitive concept of a material world. There is no agreement on a scientific definition of consciousness either, and not even a consensus that the word “consciousness” refers to anything which actually exists. And yet the *function* of consciousness should not be a mystery, because whether you ask a neuroscientist what brains are for, or you introspectively ask yourself what your own consciousness spends most of its time actually doing, you should arrive at one and the same conclusion: brains and minds are both engaged in a process of modelling a reality beyond what appears to be ourselves, with those selves as an essential part of the model, as coherent entities which persist through time. This is a process of making predictions about physically possible futures, and assigning value to the various options, in order to select a single preference. Almost all of our waking activity is devoted to two general tasks:

- (1) Trying to understand the world – both its deterministic processes and the behaviour of the other conscious beings with which we share it.
- (2) Trying to coax one potential future into being at the expense of all the others, or the reverse: trying to avoid certain possible futures. This process is goal-oriented and meaning-laden.

This description of the purpose of consciousness applies just as well to the control of the bodily movements of the most primitive conscious animals to a modern human trying to understand what is really going on in the world so they can coherently yearn for something better than the catastrophe that currently plays out. We are the ultimate biological model-makers, and we have

more information about the nature of reality than most previous generations dreamed was even possible. You might think, therefore, that we should be closing in on a coherent model of the whole of reality. In fact, nothing could be further from the truth: we have arrived some sort of nadir in this respect. We have become accustomed to living within a world whose structure we are unable to grasp. We've gathered a vast collection of puzzle pieces, some of which fit together with some other pieces, others of which don't seem to belong anywhere at all. What we conspicuously *don't* have is the picture on the box, and that makes it tricky to tell which of the candidate pieces actually belong to the real puzzle, and which are distractions, tempting us away from the only complete solution towards partial pictures which necessarily leave things out.

Scientific materialism, by definition, makes a claim about reality as a whole. It doesn't just say that the objects of scientific study actually exist – that science aims at true knowledge of objective reality. That is the more cautious claim of scientific realism, which leaves open the possibility of aspects to reality which are beyond the reach of science as we currently understand it, and might remain that way forever. Materialism claims that a material cosmos is *everything that exists*. This map of reality reached a zenith of explanatory power at the end of the 19th century, and has been declining ever since. Its successor is physicalism, which also makes a claim about the whole of reality: that it consists of whatever sorts of ontological primitives ("basic stuff") our best physical theories describe. Unfortunately, our best physics is quantum mechanics and there is no consensus whatsoever about what *that* can tell us about the structure of reality. As things stand, at least twelve major "interpretations" make rival claims about what reality is made of, how causality works, and what it is theoretically possible know about such things. These interpretations aren't scientific (that's what makes them interpretations). They are *philosophical* theories, each of which is consistent with empirical evidence (at least according to their proponents). The most significant progress we've made in the last century wasn't empirical either, but a logical proof of a kind known as a "no-go theorem", discovered by physicist John Bell. Bell's theorem rules some important things out, but raises more questions than it answers. There's been a steady stream of new interpretations ever since, and for every version of physicalism we could base on a metaphysical interpretation of QM, there's probably another two based on something else entirely. Unsurprisingly, given their abundance, no version of physicalism commands a consensus, even among physicalists themselves. It follows that there's no compelling reason to believe that *any* of them are both correct and complete. Physicalism is not a single thing but an extensive *menu* of models of reality, and people are entirely free to choose between them. Many on that side of the debate fall into scientism: the outright rejection of philosophy, with naturalistic assumptions and scientific thinking unapologetically applied to absolutely everything.

These observations aren't new. While some kind of physicalism remains the de-facto dominant paradigm for mainstream science, there is no shortage of people pointing out its deficiencies. However, that is pretty much where agreement among anti-physicalists (or anti-reductionists) ends. And since even comprehensively broken paradigms do not finally roll over and die until an alternative emerges which has sufficient explanatory power to assemble a new consensus, and since there is currently no clear sign of this happening, physicalism remains stubbornly alive. The anti-physicalist camp has a menu of its own, and it is even more exotic. Some of these systems are new, others are rehashes of ancient ideas, and one (animism) even predates civilisation itself. Some of are fully consistent with empirical evidence, while others treat physicalism's difficulties as an implicit failure of science itself, and consequently take certain empirical claims with a very large pinch of salt.

This sort of situation might be acceptable to postmodern relativists, but for scientific realists it is a serious and deepening crisis. If it is the job of science to deliver truths about reality, then why do we find ourselves so far away from a consensus about its top level structure and fundamental nature? If there's a new paradigm forming in the cracks of the old one then it is currently not sufficiently well-formed to be recognisable as such. If we assume that postmodernism and other

forms of scientific anti-realism are wrong (i.e. if we assume the puzzle has a single correct solution, and that humans are theoretically capable of finding it), this suggests that at least one critical piece is missing, or that something, somewhere needs to be turned upside-down, inside-out or back-to-front.

The problems with our current paradigm(s) cluster in three main areas. The first of these is the interpretation of quantum mechanics. Central to this is what is known as “the Measurement Problem” (MP): how does a probabilistic range of predictions about future observations become a single observed outcome? The second is consciousness. Why does it exist? What is it? What does it do? How did it evolve? What holds it together as one? Why is it that animals with tiny brains can effortlessly solve certain kinds of problems that even our best computers find immensely challenging? Why does it feel like we've got free will if our behaviour is fully determined by the laws of physics, just like the rest of the universe? The key problem here is known as the “Hard Problem of Consciousness”: If physicalism is true, then why aren't we zombies?

The third problem area is cosmology, and in this case we have a broad selection of anomalies and discrepancies, and it is deeply unclear which of them are related to which others. There is no single standout problem like the MP or the HPC, though there are several which are potentially as serious, and in this case at least some of the problems present as explicitly mathematical. In other words, our best cosmological theories cannot account for the increasingly accurate empirical data. Instead, they require us to posit vast amounts of unknown "Dark Matter" and "Dark Energy" to work, and even this sort of jiggery-pokery cannot resolve the "Hubble Tension" between recent and early universe "measurements" of the expansion rate of the cosmos. We are left asking whether the universe is composed mostly of invisible, unknown stuff, or whether the entire model, propped up by ad hoc mechanisms like inflation, is fundamentally broken.

One of the key points I will make in this book is that these three areas of problems are deeply inter-related – another proposal which is resisted by many scientists, and sometimes rather fiercely. What if these three problem clusters – quantum measurement, consciousness, and cosmology – are not three separate mysteries at all, but three different manifestations of a single, deeper confusion about the nature of reality itself? What if the reason none of them can be solved in isolation is that they all originate from the same hidden fault-line in our metaphysical foundations? What is the true relationship between mind and matter, subject and object, observer and observed? For more than a century, physics, philosophy, and neuroscience have orbited this same impenetrable mystery without recognising it as a shared problem. Each discipline encounters the same boundary where explanation fails: the point at which reality and appearance cannot be cleanly separated. It is precisely here, at this common boundary, that a new synthesis must begin.

This book is about a paradigm shift that has been brewing since Darwin removed *purpose* from nature, and then Nietzsche revealed that once purpose is gone from nature, it cannot survive in morality either. Darwin destabilised the cosmic order; Nietzsche completed the job by destabilising the moral order. Together they mark the birth of modern nihilism and the search for new foundations: the crisis out of which modernism, existentialism, and later process philosophies emerged. This paradigm shift involves the regrounding of the whole corpus of empirical scientific knowledge from the rotting foundation on which it currently sits to a new foundation which can account not just for material world we observe, but the observer itself and the underlying quantum realm which we can never observe, but which must nevertheless be part of reality.

This paradigm shift is primarily philosophical, not scientific. The main direct and immediate consequences for science itself are the dissolution or reframing of extremely intractable problems which are currently wrongly believed to be purely scientific. So here is my first empirical prediction: until they have figured out how to fix their philosophical problems, no cosmologist will be able to get their numbers to add up in an elegant and parsimonious manner. The problem, at its core, is deeply entrenched dualistic thinking, **including the belief that materialism or idealism offer a legitimate escape from dualism.** From the new perspective, each of them is revealed to be

one half of Descartes' dualism with the other crudely chopped off. Materialism is like Yang without Yin, idealism is like Yin without Yang, and together they re-enact dualism as a perennial philosophical stalemate. In reality, Yin and Yang only make sense as complementary halves of a deeper unity. The escape from dualism must be *neutral*.

Chapter 2: The whole elephant

There is an old parable, originally from the Indian subcontinent, but adopted by many other cultures. It goes like this:

Six blind men came upon an elephant in the dimness of their unseeing world. Each reached out with curious hands, eager to grasp the nature of this strange beast. One, stroking its broad, solid flank, declared: “*It is a wall, firm and unyielding.*” Another, feeling the smooth curve of a tusk, countered: “*No, it is a spear, sharp and deadly.*” A third, clasping the restless trunk, cried: “*It is a serpent, coiling and alive.*” The fourth, brushing the great ear, smiled: “*Surely it is a fan, wide and gentle.*” The fifth, embracing a sturdy leg, proclaimed: “*It is a tree, rooted in the earth.*” The last, tugging at the tail, laughed: “*It is a rope, rough and dangling.*” So they quarrelled without end, each certain of his truth, each blind to the others’.

In this book, the elephant is reality itself. All previous societies tried to describe it, many believed they had succeeded, and some of them had a firmer grasp on crucial parts of the puzzle than we do. What no previous society had was modern science and technology, so their understanding of reality can only have been based on a mixture of experience, reason, faith and revelation. We are now in possession of a vast amount of reliable information about how reality works, all of which must be taken into account. Unfortunately, science doesn't describe the whole elephant either – it is nowhere near providing a coherent, integrated model of reality, even by its own limited illumination. And while neither faith nor revelation can contribute anything useful to a new foundational paradigm, both reason and experience are indispensable, although neither of them qualify as science. Just because an item of knowledge can be acquired before any science begins, or can only be obtained subjectively if it can be obtained at all, it doesn't follow that it is as irrelevant as faith and revelation. We don't need science to tell us we are conscious, for example. In fact, science *can't* tell us this. If we want to know if someone is conscious, we must ask them and evaluate their answer.

Science isn't just one of the blind men but several. Cosmology isn't integrated with quantum mechanics, and neither of them make firm contact with any theory of consciousness. Even *within* these great areas of academic study there is profound disagreement about central questions. Science doesn't know what consciousness is, and doesn't know what quantum mechanics means, and our best cosmological model, which goes by the not-so-snappy name of Lambda Cold Dark Matter (Λ CDM), has got more holes in it than a Swiss cheese. There is deep confusion about the expansion rate of the universe, the cosmological constant problem is the biggest discrepancy in the history of physics, we don't understand what gravity is, we can't adequately explain why the early universe was so exceptionally uniform or why the constants are fine-tuned for life, and while our theories suggest that life should be abundant in the cosmos, we can find no sign of it anywhere but Earth. However, what really damns Λ CDM is the direction of travel. It becomes ever more precise (frequently at vast public expense) but instead of leading to solutions to the biggest problems it regularly reveals new ones. Meanwhile, the discrepancies we're already aware of are being measured to ever greater resolutions, as if higher precision could somehow compensate for a lack of fundamental insight. As things stand, Λ CDM cannot even be relied on as an accurate account of just one of the elephant's body parts, let alone the whole thing. Something is fundamentally wrong with it, and cosmologists don't know what. Almost all of their new ideas are proposals for yet more ad-hoc additions to Λ CDM, rather than something that could replace it wholesale. In other words, we

are justified in believing Λ CDM is the best materialistic theory available, but as things stand, it is **empirically inadequate**. Very literally, today's cosmologists cannot find an elegant and parsimonious way to make their numbers add up, and in some situations they can't make them add up at all.

The other blind men, of course, are philosophy and religion. Western philosophy, split into its Analytic and Continental traditions, has long grappled with the elephant by way of reason alone. Analytic philosophers try to tame it with language and logic, dissecting concepts as if the creature could be understood by cataloguing the texture of its skin or the angle of its bones. Meanwhile the Continental philosophers, after much struggling, eventually threw up their hands and declared that the very idea of a single, unified model of reality is not just a wild goose chase but a despicable form of oppression. They tell us we must just resign ourselves to epistemic chaos, and celebrate the diversity of “mini-narratives”, even if they directly contradict each other or don't even make sense on their own.

Which naturally brings us to religion. The Abrahamic traditions gave us stories of creation, covenant, and salvation, and declared the elephant to be the finest possible work a transcendent lawgiver. Eastern mystical traditions – Hinduism, Buddhism, Taoism – sought instead to dissolve the boundaries of perception, to experience the elephant from within, as an inseparable flow of being and consciousness. These perspectives uncovered profound truths, but as things stand, there is no successful integration of these insights with our scientific models of the structure of reality.

Few people have even attempted a full synthesis of science and spirituality, though there have been some notable exceptions. Ken Wilber is the best known living example, and if I'm looking for recent historic examples then I'd choose three who articulated visions as the quantum age was being born and three from the later 20th century. The three earlier thinkers were responding to the same predicament that QM itself exposed: the breakdown of mechanistic materialism. All three lived *through* the initial quantum revolution. They can be seen as the first wave of integral cosmologists anticipating what the quantum paradigm would eventually demand: a participatory, processual, value-laden universe.

Alfred North Whitehead (1861-1947) was a philosopher of process thought, who saw reality as a "web of becoming" where science, philosophy, and religion interpenetrate. *Science and the Modern World* (1925) was written as quantum theory was coalescing, and *Process and Reality* (1929) came two years after the Solvay Conference where Bohr and Einstein debated the meaning of quantum indeterminacy. Many modern process philosophers interpret Whitehead's thought as a philosophical parallel to quantum ontology, even though he developed it independently.

Sri Aurobindo (1872-1950) was Indian philosopher-yogi who sought a synthesis of Eastern spirituality with Western evolutionary thought. His main metaphysical synthesis was underway before QM's formal development but was revised during 1930s. Aurobindo's ontology of reality as a dynamic manifestation of consciousness, with matter as involved spirit, mirrors quantum indeterminacy's dissolution of fixed matter. He was spiritually articulating the same shift from substance to process and from object to participation that physics was uncovering empirically.

Pierre Teilhard de Chardin (1881–1955) was more influenced by Darwinian and thermodynamic cosmology than by QM. He developed his evolutionary cosmology during the 1930s–40s, and like the quantum pioneers, he saw consciousness and complexity as integral to the fabric of evolution. His cosmogenesis anticipated the idea of a participatory universe that later physicists would develop explicitly under quantum influence.

The later three each deepened the integration of scientific and spiritual cosmology in ways that directly reflected the maturing quantum and ecological worldviews of their time. If the first three thinkers anticipated the quantum shift philosophically, this later trio sought to live within it – to think from *inside* a participatory cosmos rather than about it from the outside.

Thomas Berry (1914–2009), a cultural historian and eco-theologian, gave this participatory cosmology its most elegant expression in the phrase, “*the universe is a communion of subjects, not*

a collection of objects.” Berry’s vision placed the human story within the larger unfolding of the cosmos, framing the ecological crisis as a crisis of cosmology – a breakdown in our sense of participation in the great story of creation. He called for a new cosmological narrative, inspired by the emerging insights of systems theory, deep ecology, and quantum interconnectedness. For Berry, science and spirituality are not separate enterprises but two modes of intimacy with the same sacred reality.

David Bohm (1917–1992), a theoretical physicist and close associate of Einstein, provided the most rigorous scientific articulation of this participatory worldview. His theory of the *implicate order* proposed that the manifest world (the *explicate order*) unfolds from a deeper, undivided wholeness. In this hidden dimension, every part “enfolds” the whole, and consciousness participates in the “unfolding” process itself. Bohm’s dialogues with Jiddu Krishnamurti and his later work on thought, meaning, and creativity show how far he extended physics into the domain of metaphysics and spirituality. He attempted to dissolve the boundaries between subject and object, offering a scientific expression of the same participatory ontology that others on this list approached through philosophy and theology. However, Bohm’s system is entirely deterministic, and I believe it evades, rather than solves, the measurement problem. I shall have more to say about this in chapter 6.

Henryk Skolimowski (1930–2018) was a Polish philosopher of eco-cosmology who developed what he called an “eco-philosophy of participatory mind.” Trained in both science and philosophy, he argued that the universe is not a mindless mechanism but a meaningful, evolving field of consciousness in which human awareness participates creatively. He saw the ecological crisis as a spiritual and epistemological failure: the direct result of viewing the world as dead matter. His alternative was a sacred cosmology grounded in reverence for life and participation in the cosmic process of meaning-making. In this sense, he extended the implications of quantum indeterminacy and observer participation into the moral and ecological domain, proposing that the universe itself is a value-realising system.

Together, these three thinkers represent a second wave of integrative cosmology, which moved beyond mere synthesis toward a new worldview in which science itself becomes a spiritual practice in itself. Each, in his own way, recognised that the participatory universe implied by quantum theory and ecology is not only a description of reality but a call to responsibility: to awaken, to participate, and to co-create the next phase of the cosmic story.

There is one other person I cannot leave unmentioned, because he was such an important influence in my own case. Aldous Huxley (1894–1963) served as the great communicator of this integrative turn. Through works like *The Perennial Philosophy* (1945) and *Island* (1962), he distilled the insights of both Eastern mysticism and Western science into a vision of the universe as mind in evolution. Though less a cosmological system-builder than the others I have mentioned here, Huxley did more than anybody else to transmit the participatory worldview to modern culture, making the cosmic synthesis accessible to the broader imagination.

The ideas explored in this book are related to the thinking of all of these individuals, as well as some more recent thinkers whose work focused on getting specific parts of the picture more into focus – most notably philosopher Thomas Nagel and physicist Henry Stapp. Between them, Stapp’s *Mindful Universe: Quantum Mechanics and the Participating Observer* (2007) and Nagel’s *Mind and Cosmos: Why the Materialistic Neo-Darwinian Conception of Nature is Almost Certainly False* (2012) provided me with a new vision of the picture of on the box. On their own, neither would have been enough.

The whole elephant

I am an ex-materialist. Once upon a time, in what now seems like a previous life, I was the administrator for the newly-created bulletin board of the Richard Dawkins Foundation. Around the same time, I arrived at the conclusion that materialism doesn't actually make any sense, because the only way to make it internally coherent is to deny that the word “consciousness” refers to anything real. This is known as “eliminative materialism” (or just “eliminativism”) and its logical consistency does nothing to diminish its absurdity: the truth is that it is only through consciousness that we are aware that *anything* exists at all. That was back in 2002, and I have spent much of the intervening period both exploring what a coherent post-materialistic model of reality might actually look like, and trying to find ways to prise open the minds of people who think like I did until my “defection to the light side” at the age of 33. The first activity has proved very rewarding...eventually; 25 years later I am ready to tell a new story. The second turned out to be almost impossible. I found a way to *break out* of a materialistic belief system, but nobody could have *broken in* and rescued me from it, not least because I did not feel I was in any need of rescue.

This raises an obvious question. If materialism can be falsified with a combination of pure reason and an acknowledgement of the reality of consciousness then why has it retained its position as the dominant metaphysical ideology of modernity? Why hasn't it been displaced by a new paradigm? On one level the answer is simple: there is no coherent new paradigm to displace it. Materialists typically frame the situation as a straight choice between materialism (which they assume to be an unproblematic default starting premise) and dualism (which is what you inevitably end up with if you add something to materialism). Meanwhile, very few people who reject materialism actually claim to be (or should I say “admit to being”) a dualist. Some call themselves “non-dualists” in order to ram this point home, although that term has many different meanings. The clear positions opposed to materialism could be categorised into three main groups: idealists (who believe consciousness is everything), panpsychists (who believe everything is conscious), and “other / don't know” (mostly people who understand what is wrong with materialism, but aren't convinced idealism or panpsychism are true either, usually because they consider brains to be *necessary* for consciousness *even if insufficient*). To materialists, all of it looks like woo, and because there are multiple incompatible alternatives on offer, nothing changes. It's a stalemate: unresolvable metaphysical trench warfare.

That said, there are quite a few *parts* of this new paradigm coming into focus. Based on the current state of books written on this topic the “whole elephant” should look something like this:

- **Reality is not fundamentally material but relational and experiential.** Matter, mind, and meaning are not separate domains but aspects of a deeper unity.
- **Consciousness is not an anomaly but a principle woven into the fabric of the cosmos.** It is as basic as mass, energy, or spacetime, and perhaps more so.
- **The cosmos is participatory.** Observation, valuation, and relationship help shape what is real, not just passively register it.
- **Time and process are fundamental.** Being is not a static block but an unfolding, in which novelty, emergence, and irreducible subjectivity matter.
- **Ecology and interconnection are the true grammar of existence.** From fungi and forests, brains to quantum events, the world is a web of mutual becoming, not a collection of separate objects.
- **Meaning and value are ontological, not epiphenomenal.** They belong to the structure of reality, not just to human projections.

In one sentence the *missing paradigm* is a participatory, meaning-infused, relational cosmology where mind, matter, time, and life are continuous aspects of one living process: the universe as a communion of subjects.

This is a pretty good start. But if we can get this far, why can't we find a way to agree on the details to a sufficient extent that a coherent new paradigm can emerge, and begin the process of displacing materialism? Is it simply because not enough people have got the message? I don't think so. I think that if a proposed new paradigm actually had enough inter-disciplinary explanatory power, then the paradigms would already be shifting. Something must therefore be missing, and whatever it is has got to be as elegant and conceptually simple as shifting the centre of the Solar System from Earth to the Sun. Is it possible that there is some relatively simple way of re-arranging the puzzle pieces so that everything makes sense in a radically new way? And if so, why haven't we already figured out what it is? Why can't we just get an AI to analyse the whole situation (at which it is apparently so astonishingly effective), and come up with the correct answer?

My answer is this: if we accept that reality is coherent, consistent and comprehensible then it is not possible for the idealists, panpsychists and "other / don't know's" all to be correct. If there actually *is* a whole elephant then two of those groups will necessarily have to join the materialists in accepting that some of their foundational beliefs need to change. In fact, this sort of situation applies to just about every Western mind currently in existence: as things stand there is no serious candidate for the right answer, so each of us is free to believe whatever we like, and we can justify this at no cost by saying, and fully meaning, that we believe it to be the least bad option on the table. Collectively, this situation suits us. We *like* it. Sure, it might lead to arguments and meaninglessness, but it is the ultimate manifestation of an individualistic theme that has always run through Western societies. At this point I start to wonder who is left that will still want to consider something radical enough to power a paradigm shift of the magnitude required. Is the loss of a wide variety of incoherent worldviews a price worth paying for a coherent model of the whole of reality?

Most people do not understand the severity of this situation, and therefore neither accept it as normal nor view it as a problem. The fundamental fragmentation – the incoherence between science, philosophy, and spirituality regarding the nature of reality – is primarily an internal crisis debated within academic and intellectual circles. Most people live with an implicit, common-sense materialism (the world is made of stuff, and my mind is in my head) that is never challenged or explicitly articulated. As for the cosmological crisis, while the public enjoys popular science books about the Big Bang and Dark Matter, they generally absorb Λ CDM as a settled scientific fact, not a structure riddled with fundamental, widening cracks. The consensus is presented, the confusion is not. The problems with QM *have* seeped into public consciousness, but often in a distorted, pseudoscientific fashion. People know about "quantum weirdness" and superposition, and have heard about the controversial role of the observer, but this is often misinterpreted as proof of whatever belief they already hold, rather than understood to be the precise, century-old conceptual failure of the dominant physical theory that it actually is.

The fragmentation is therefore normalised. The most relevant public awareness is not of the academic theoretical problems, but of the consequences for society in general: the lack of a shared, coherent story for living. The average person *expects* to choose from a variety of belief systems. The idea of a single, coherent model that integrates all these domains is not a common expectation. Meaninglessness is also accepted. Ours is an era of epistemic relativism ("my truth" vs. "your truth") where people think it is normal that collective, objective meaning is impossible. This nihilism, paradoxically, has also become a normalised part of the modern condition. The intellectual stalemate is felt as a kind of cultural and spiritual drift, but the cause is unknown. A coherent model of reality would demand intellectual honesty and a willingness to surrender cherished fragments of supposed truth that no longer fit any potential whole. That's a tall order. Many people (spiritual seekers, materialists and postmodernists alike) have built their identities around their preferred paradoxes. These contradictions aren't just tolerated; they're celebrated as signs of depth or

sophistication. Incoherence offers freedom from commitment (you can dabble in ideas without being bound by them), immunity to critique (if your worldview is self-consciously paradoxical, it can't be "wrong") and aesthetic appeal (contradictions feel poetic, mysterious, even sacred). Coherence threatens all that. It threatens the buffet of metaphysical options that lets each person pick their own flavour of reality. Except, ironically, the freedom to believe in nonsense is itself a kind of tyranny. It problematises shared understanding, undermines trust in reason, and makes collective meaning-making nearly impossible. We end up with epistemic relativism, cultural nihilism and philosophical gridlock: no paradigm shift can occur because every worldview is protected by the general epistemic anarchy. In this light, coherence is liberation. It's the promise of a worldview that *works*, that *adds up*, that *honours experience without collapsing into chaos*. What we actually fear is the loss of our metaphysical playground, and I think that fear is misplaced. A coherent model of reality doesn't have to be rigid or reductive, and it certainly doesn't have to be authoritarian. What it can't be is endlessly self-contradictory. So is the loss of all those incoherent worldviews worth it? I'd say: only if the new paradigm is rich enough to absorb their insights without their contradictions. The challenge is replacing the menu with something better.

There *is* a new story available, and it *does* satisfy the description of the whole elephant given above, but it does so in a way that almost nobody is expecting. And if you think about it, then I hope you'll agree that it always had to be that way.

A new story

I have a new story to tell. In this story, the universe that appears to us – a three-dimensional realm containing material objects which dynamically changes over time – did not begin 14 billion years ago in a "big bang" which came from nothing, for no reason. In this story, that kind of universe can only exist if it contains at least one thing which is capable of subjective experience. This thing is a creature which doesn't just react, but *thinks*. A being which can construct a model of the world, with itself in the model, and which intuitively understands that multiple futures are physically possible, and can assign value to the various options. Something, in other words, which can *select a best possible world*, even if only from its own limited, subjective perspective.⁴

This story appears to have *two* beginnings (and neither stories nor universes are supposed to have two of those), but I am not suggesting that the 4 billion year old Earth, replete with a fully-functioning biosphere, popped into existence out of nothing, for no reason. I am suggesting it emerged from a deeper level of reality, where nothing begins or ends, there is no present moment, and time is an unchanging block of information that encodes every structural possibility. This realm is mathematical and structural, not physical or mental. It is the potentially infinite realm of formal possibility. Its ultimate foundation is Anaximander's Apeiron, which I call "Phase 0", within which is the Pythagorean realm of number, or structure, which I call "Phase 1". "Phase 2" – the aforementioned *universe that appears to us* – is the realm of actuality and becoming – of matter within mind. Here it is always now, everything has beginnings and endings, and change is ever-present as time inexorably flows, always (at least it seems) in the same direction.

Together this makes the Two-Phase Cosmology (2PC) and it is essential to note that "phase" can apply in two different senses here. The first regards the historical structure of the cosmos: two stages in the history of both cosmological and biological evolution, before and after the appearance of the first conscious animal (I call this animal LUCAS – the Last Universal Common Ancestor of Subjectivity). The second regards its metaphysical structure: two different states of reality. In historical (or "primordial") Phase 1, which is eternal and unchanging, there is no ontological Phase 2. But in historical Phase 2, which isn't primordial and by definition includes the present, there is a

⁴ It could theoretically also be some sort of quantum computing machine which replicates what brains do, but not like any of today. Also, machines can't exist unless there's a creature capable of constructing them.

kind of ontological Phase 1 ever present in the background. The difference is that now it too is dynamically changing, because it is locked in a cosmic duet with Phase 2 (so I call this “participatory Phase 1”). The process of selecting a single Phase 2 actuality from an almost infinite set of Phase 1 physical possibilities is what we call consciousness.

The pivot around which reality balances

In this model the shift from Phase 1 to Phase 2 is a dynamic pivot around which reality continually rebalances itself. You might be wondering why, if this strange story is actually an accurate model of reality, cosmologists and physicists have not suggested it before. In fact, just such a pivot is already present in our scientific model of reality – or at least it is *trying* to be. But there is a problem, and it is a humdinger. The present scientific model is inescapably physicalistic, and since there can't be two different realms or levels of reality if physicalism is to be taken seriously as a monism, there is no satisfactory place to put a pivot. This is why quantum mechanics seems so incomprehensibly weird. The Measurement Problem is the most profound internal conceptual problem in physics, and arguably the second biggest conceptual problem in the whole of science.

Many have tried to find a way to place this pivot within the material world of phase 2, but after a century of these attempts, nobody has been able to make it work. The problem is that a physical pivot of this sort, if real, must be empirically verifiable. Isn't that the whole point in assuming it is physical in the first place? If it becomes impossible to progress beyond arbitrary untestable hypotheses or pseudoscientific distinctions between a “micro world” and “macro world”⁵, then the theories containing this sort of pivot start to look suspiciously like *failed science*. In physicalism there's supposed to only be *one* world, and if you're going to say there's two then the relationship between them cannot be vague or untestable.

This ongoing failure has led another significant group of scientists and philosophers to a radical conclusion first suggested by Hugh Everett in 1957: there is no pivot. This combines my two phases into one, resulting in a realm not of infinite possibility, but infinite actuality. The singular nature of our experience, they say, is an illusion. In this model of reality, whenever it is physically possible for more than one thing to happen (which is all the time), *everything does*. And since no exception is made for anything as cosmically irrelevant as conscious beings, it follows that our minds are also continually splitting. Instead of just the one life, we're living an unimaginably huge number in parallel, and every time it feels like we are choosing between different physically possible futures, we actually make *every* possible choice. This “Many Worlds” interpretation of quantum mechanics (MWI) is popular in Hollywood and has been championed by a few brave individuals, but I doubt that many people actually believe it is true. Its status is more like “least bad option available” than “convincing model of how reality actually works”. The root problem is the strength of the subjective impression that we have free will. Most of us are convinced that we can freely choose a specific future, at least for the movement of own bodies, and that this entails *not choosing any of the others*. We are convinced that we can marry Alice or we can marry Bob, but we cannot marry both of them. MWI denies this. One is left wondering what brains are for, if the decisions they make don't make any difference to what actually happens.

This leads us to the other kind of solution to this problem that is currently on offer – and it is the most controversial of all, because it directly links the second biggest conceptual problem in science to the first. It is often assumed to have been invented relatively recently by somebody like Deepak Chopra. In fact it was first proposed in 1932 by the most influential mathematician of the 20th century, and a genuine contender for the greatest mathematician/physicist who ever lived. In his

5 Yes, this specific distinction is pseudoscience, and the worst kind, because it originates within the scientific castle walls. It is invariably touted as scientific orthodoxy, but it is supported by zero empirical evidence, which is why there's no consensus whatsoever as to whether or not is actually true.

ground-defining book *The Mathematical Foundations of Quantum Mechanics*, John von Neumann's only interest was formalising quantum theory, but he faced the problem of nobody having the faintest idea where in the emerging model of reality this pivot between possibility and actuality should go. Von Neumann pointed out that, according to the theory itself, the pivot could be placed anywhere between the quantum system being measured and the consciousness of the human observer. He also noted that consciousness is the only place in this chain of causality that is in any way special, and the only point where we can say, with total certainty, that the collapse has taken place. Von Neumann therefore left the “collapse event” out of the mathematical formalism, and suggested that consciousness is involved. Mainstream science has always been hostile to this proposal, primarily because it explicitly denies materialism. Many people also wrongly believe that it can be definitively rejected as false. The truth is that the problem which initially caused von Neumann to remove collapse from the physical system has never been solved. A century after the discovery of quantum mechanics, we are no closer to resolving the Measurement Problem than we were at the beginning.

Identifying the correct solution is the key to unlocking a new paradigm, but as things stand, there is no consensus as to what sort of answer we should be looking for, no agreement on how we might recognise it if it shows up, and not even any guarantee that a correct answer will present itself at all. Maybe our brains just aren't built to understand such things.

Schrödinger's second equation

To understand why this mysterious “collapse” matters so much, we need to look at the central concept of quantum mechanics: the *wave function*. The wave function is not a physical wave moving through space, but more like a giant map of everything that could possibly happen. If you take something small enough (say, an electron) the wave function doesn't tell you *where it is*. Instead, it tells you *all the places it could be*. Before we look, the electron isn't in one place or another. In a very real sense, it is in all of them at once.

Erwin Schrödinger was the first person to write down the equation that tells this strange “possibility wave” how to move and change. This was in 1926, and the result, known as the Schrödinger equation, is a cornerstone of modern physics. But from the beginning there was a catch: the equation describes how possibilities spread and evolve, but not how they resolve. It can tell us all the futures (or pasts) that might happen (or have happened), but it cannot explain why only one of them does (or did).

Schrödinger himself struggled with this gap. Later in his life, he turned toward philosophy in search of an answer. What he found was not in Europe's scientific tradition at all, but in India's ancient Upanishads. These texts describe a profound identity between the root of personal consciousness (*Atman*) and the ground of all being (*Brahman*). Schrödinger believed that this insight solved the riddle of collapse: the transition from possibility to actuality is tied to consciousness itself, and consciousness is not just a small private bubble inside our heads but the very ground of reality.⁶ He famously, if somewhat mischievously, referred to this as his 'second equation'. Though it never appeared in a peer-reviewed journal, it was the necessary philosophical bookend to his work in physics: an identity statement that claimed to reveal the only thing that can bring those possibilities into a single, actual world: the unity of consciousness with the universe itself: *Atman equals Brahman*. However, Schrödinger never explicitly integrated his two equations – for him they remained parallel ways of understanding reality, rather than components of an integrated metaphysical and physical system. The primary purpose of the present book is to show how this integration might actually work. Schrödinger was talking about the same pivot I have been

6 Whether Schrödinger was an idealist or a neutral monist is debatable, but his second equation does not imply idealism, and most modern translators think Advaita is closer to a pure neutral monism than to idealism.

talking about. His first equation describes what “happens” in what I call Phase 1. His second is an essential part of the explanation for what happens in Phase 2; collapse is where these two phases meet. Schrödinger’s proposal was that this pivot cannot be found in physics alone. It can only be found by recognising that the root of the consciousness which experiences reality from within is the same thing that brings reality forth at all.

This was too radical for physics to absorb. The first equation could be tested, measured, and turned into technology. The second was a metaphysical claim about the ground of being. Yet Schrödinger saw them as two halves of the same deeper truth.

Although the origin of this hypothesis is mystical (to the extent that it could serve as the broadest possible definition of what the word “mystical” actually means), what follows is not a religious narrative. Things I am **not** implying include: God⁷, idealism, disembodied minds of any sort, individuated souls and the afterlife, revealed moral rules, or anything to do with spiritual development. I am attempting to do philosophy. I will never be any sort of guru and the purpose of this book is not “self-help”. The hypothesis is that the root of personal consciousness – the observer rather than what is observed – is ontologically identical to what I formally write as $0|\infty$, and call “the Void”. **In this book I shall be treating Schrödinger's second equation not as a mystical declaration, but as a *structural hypothesis*.**

7 At least according to the definitions of most Westerners.

Chapter 3: The broken paradigm

I have repeated it like a broken record: we lack a coherent model of reality. We have more data and durable theories than ever, but we are nowhere near assembling them into a coherent whole. This applies not just to science but to the whole of Western philosophy, which is divided into two "traditions" which barely acknowledge each other's existence, and which have a radically different response to the incoherence of the current situation. Continental philosophy has declared postmodern relativism to be the last word, and not only given up on the search for a big picture, but declared the very idea to be an authoritarian anachronism. Analytical philosophy is still committed to that search, but there is zero consensus with respect to which direction progress lies. Only theology provides a unifying answer, but it does so by attributing the cosmos to God's will and declaring that ultimate understanding is beyond the imperfect minds of humans.

We can do better than this. Surely reality must have a specific structure, and I can see no reason to believe that humans are incapable of discovering it and understanding it. Maybe not understanding *everything* about it – some things may necessarily remain mysterious forever – but understanding enough about it to rid ourselves of the paradoxes, discrepancies, anomalies and unanswerable questions which currently plague cosmology, quantum mechanics and cognitive science (which is the closest thing we've got to a science of consciousness).

Key definitions

This chapter provides an overview of the problem space, and I must start by defining the most important technical terms. I am going to argue that these problems are ultimately *conceptual* in nature – that they are the result of a misalignment between our metaphysical vocabulary and the actual structure of reality. Definitions like this are therefore of major importance, which inevitably results in serious disagreements about what these words *should* mean. There is no point in attempting to describe this conceptual battlefield in terms of the disagreements about definitions; I'd never get past the definitions and nobody would want to read the resulting book. **The definitions I give here apply throughout the remainder of this book. They only need to be reasonable and coherent for the purposes of my argument, even if others use the same words differently.**

The perfect example of this terminological confusion are the words "material" and "physical". Many people use them interchangeably, and don't acknowledge any important difference between "**materialism**" (the belief that reality is made of material objects, or that the material universe is all that exists) and "physicalism" (ditto). Unfortunately (for physicalism) we have two radically different concepts of "physical" – those which came before and after the discovery of quantum mechanics. "**Physicalism**" is typically defined as "the belief that reality is made of whatever our current best physical theories indicate that it is made of", which would be great if scientists (or philosophers) could agree about what quantum mechanics is telling us about the nature of reality. Unfortunately, they can't. "Material" is less problematic as a *term* only because it more obviously refers to the pre-quantum concepts of physical, but if that is how the term "materialism" is to be defined, then as an ontological *theory* it can be certified dead.

The word "consciousness" suffers from similar problems. As things stand there is no agreement on how to define this word, and consequently no agreement about how it relates to the rest of reality, or even whether it even exists. This problem is directly related to the confusion surrounding "material" and "physical", and it could all be cleared up as follows.

"**Consciousness**" is the only reason we know reality exists at all. It is the frame for our own

subjective experience of reality. As such, the only way we can define it is in terms of subjectivity itself – we must, in effect, mentally point to our own experiences and associate the word with those experiences. The technical name for this is a "private ostensive definition". This is not an orthodox definition, but it establishes what the word is supposed to mean. I must stress that Wittgenstein's private language argument does not apply in this case, because I am not trying to define a private language; I am merely establishing the meaning of the word "consciousness" *as used*. We can take this a bit further, because it is necessary to ensure that we avoid **solipsism** (the belief that nothing exists outside our own mind). Within our own consciousness, we are aware of a large number of other beings which behave as if they are conscious – not just other humans but also most animals, right down to the level of something like an insect or a worm (although how where exactly we draw the line is very much an open question at this point). If we assume these other beings are actually conscious too, then solipsism can be dismissed.

We can now give clear definitions of material and physical. The "**material world**" is a three-dimensional realm, populated with objects both living and non-living, which continually changes as time passes. In this world it is always the present moment, time always "flows" in the same direction, and objects are always in just one place at any one time and have a single set of properties. We are intimately familiar with this material world, because it is presented to us within consciousness whenever we're awake. When we're asleep we experience a "ghostly" version of the material world – one which seems real enough to the dreamer, but events which occur within it are not constrained by the laws of physics.

The "**physical world**" is something that forever lies beyond the veil of perception – we can never escape our own consciousness and experience that world directly. Not everybody agrees that this mind-external world should be called "physical" – physicalists and dualists do, but **objective idealists** claim it is another sort of consciousness and **subjective idealists** deny that it exists at all. I think we must assume that it *does* exist, for there must be a reason why things remain consistent in each of our individual versions (or "projections") of material reality. I shall use "physical" to refer to this objective reality (and I will explain why), but I must make clear that I'm not stipulating that the parts of it that correspond to the material world are all that there is – I will leave open the possibility that other things might also exist in the objective realm beyond our minds, because I don't think we are in a position to safely rule them out (though skepticism is reasonable). In this way I can use "physical" to refer to the parts that *do* correspond to our projected material realities. This mind-external realm can also be called "quantum reality", and QM implies that objects can be in multiple places and have multiple sets of properties until such time as they are measured or observed (whatever *that* means). It is also far from clear whether there is any such thing as "now" in this realm or whether time flows forwards or backwards (or perhaps a bit of both, or neither). In other words, the concept of the physical world I am describing now is a very different to concept to that of the material world with which we're so familiar. This justifies the use of two different terms.

The reason we can't agree on whether this objective part of reality should be labelled mental, physical or something else is because all we can actually know is its structure – or part of its structure. We can therefore clearly say this is a form of **structural realism**. But as all we can know is structure, I describe it as *neutral*. (So to be clear: I use "physical" to refer to the objective correlate of the material world, and "neutral" to describe the entirety of objective reality, which may or may not consist entirely of the physical world). The model of reality that I'm going to describe is indeed a new form of **neutral monism** – the belief that objective reality is made of neither consciousness nor matter, but something else. The only sensible candidate for the neutral stuff is *information*, though this leaves unanswered the question of what, if anything, that information is grounded in. Where did it come from? How is it stored and updated? Neutral monists have been rare in the history of western philosophy, although there are a few well known examples (Baruch Spinoza, Ernst Mach, Alfred North Whitehead, Bertrand Russell and David Bohm among them, though each of believed something different about the identity of the "neutral stuff").

An introduction to the problem space

With the initial definitions out of the way, I can now provide brief descriptions of the problems I have been talking about. I will assume readers have no knowledge of any of them, and most can't be adequately explained in one short paragraph. I will be examining all of them in greater detail in Part Two. My purpose right now is just to give an idea of their diversity and severity. Please note that some of these problems overlap. They cluster into three areas (cosmology, quantum mechanics and consciousness) but the boundaries are messy, some fall into more than one of these areas. One or two don't sit comfortably anywhere.

Cosmology

Λ CDM is every bit as broken as Ptolemaic geocentrism was in the 16th century. It consists of an ever-expanding conglomeration of ad-hoc fixes, most of which introduce as many problems as they solve. The following list may seem sprawling, but that is indicative of the intractability of the underlying situation. These problems cannot be cleanly classified because cosmology itself has no unified theory that can make sense of them. Instead, each anomaly is patched in isolation, creating an overall model that is riddled with contradictions.

1) How can something come from nothing?

There are countless ways of restating this question. Why does anything exist? Why isn't there just nothing? What caused the Big Bang? etc...

2) The Constants Fine-Tuning Problem

The fundamental constants of nature appear to be exquisitely calibrated to allow for the existence of life. Why does the universe appear to be precisely set up to make life possible?

3) The Low-Entropy Initial Condition

The universe began in an extraordinarily smooth, low-entropy (highly ordered) state, as shown by the near-uniform cosmic microwave background (CMB). Physics does not demand or explain such fine-tuning.

4) Inflation-related fine-tuning problems

To address problem (3) above and also problem (6) below, cosmologists proposed “inflation” – a fleeting period of superluminal expansion that smoothed the early cosmos. Inflation ends when its driving potential energy decays into matter and radiation, a process called reheating. For today's universe to emerge, this reheating must occur with extreme precision in both timing and efficiency, yet no known mechanism explains this. Inflation therefore fails to avoid fine-tuning, because it actually requires more fine-tuning than it gets rid of.

5) Other fine-tuning problems.

Countless additional fine-tuning issues exist. The universe shows an unusually favourable balance of elemental abundances for stable stars and biochemistry. Galaxies and stars also formed at just the right time – early enough for life to evolve, but not so early as to disrupt cosmic smoothness. Further tunings include the matter–radiation equality and primordial perturbation amplitude problems.

6) The Missing Monopoles

Grand Unified Theories (GUTs) of particle physics predict the production of magnetic monopoles – massive, stable particles carrying a net magnetic charge – during symmetry-breaking transitions in the early universe. The problem is that no magnetic monopoles have ever been observed. Inflation solves it by “diluting” them with empty space.

7) The Baryon Asymmetry Problem

A foundational assumption of particle physics and cosmology is that the laws of nature are nearly symmetric between matter and antimatter. In the earliest moments after the Big Bang, the universe should have produced equal quantities of baryons (matter) and antibaryons (antimatter) through high-energy particle interactions. What we actually observe is a universe composed almost entirely of matter.

8) The Hubble Tension

This is a large and persistent discrepancy between two different (early universe vs recent) measurements of the rate of cosmic expansion.

9) The S8 Tension

This is a persistent mismatch between the level of matter clumpiness predicted by Λ CDM for the early universe and what we actually observe in the late universe. CMB measurements fix a precise value for how strongly structures should have grown, but weak lensing, galaxy clustering, and cluster counts all find a smoother cosmos with a significantly lower S8. The gap has widened as data improved, creating a second major early-versus-late tension that the standard model cannot resolve.

10) "Dark Energy"

Dark energy was invented to account for a surprising set of astronomical observations that contradicted long-standing expectations. A repulsive force appears to be *pushing the universe apart* at an accelerating rate (almost like anti-gravity). Today, Dark Energy accounts for roughly 70% of the total energy density in standard Λ CDM, but its origin, nature, and ontological status are unknown.

11) The Cosmological Constant Problem

Dubbed "worst theoretical prediction in the history of physics", the cosmological constant problem is a staggering mismatch between theoretical prediction of the repulsive force described above and the observational measurement of that force. The mismatch is between *60 and 120 orders of magnitude*.

12) "Dark Matter"

Dark Matter has never been directly detected, but regardless of that it is now thought to comprise approximately 85% of the matter content of the universe. The hypothesis of Dark Matter emerged as a unifying explanation for multiple independent observational anomalies across different astrophysical and cosmological scales. In each case, visible (baryonic) matter alone proved insufficient to account for the observed gravitational effects. After decades of experiments, we still have no clear idea what it is or where it came from.

13) The Quantum Gravity problem

A central goal of theoretical physics for nearly a century has been the unification of quantum mechanics and General Relativity, but the two most successful theoretical frameworks remain conceptually incompatible.

14) The Early Galaxy Formation Problem

The James Webb Space Telescope has detected massive, metal-rich, well-formed galaxies at redshifts greater than 13 – meaning they already existed 325 million years after the Big Bang. The abundance, size, composition and apparent maturity of these early galaxies outpace the predictions of hierarchical structure formation, challenging both the timeline and mechanisms of Λ CDM.

15) The Fermi Paradox

Our theories suggest life should be abundant in the cosmos, but after over a century of intense searching, we have found no sign of it. Where *is* everybody?

16) The Black Hole Information Paradox.

The black hole information problem asks whether information that falls into a black hole is lost when the black hole evaporates via Hawking radiation. Modern approaches suggest that unitarity is preserved, but only by abandoning naïve locality, independent interior–exterior descriptions, or observer-independent global states. This raises a deeper conceptual question: what counts as information, where does it reside, and when does it become physically real?

17) The Arrow of Time and the Problem of Now

Human experience and natural processes clearly distinguish past from future, yet the fundamental laws of physics are time-symmetric, treating both directions equally. Why, then, do we perceive a one-way arrow of time? A related puzzle concerns the present moment: in relativity, time is just another dimension, and all events coexist in a four-dimensional block universe with no privileged “now.” Yet the present is all we ever experience.

18) The Memory Stabilisation Problem

Though rarely noted, this issue is fundamental. Memory underpins continuity, identity, and meaning, seeming to refer to fixed past events encoded as stable traces in the brain. Yet in a quantum universe where events become definite only upon observation, it remains unclear how the apparent solidity of the past, and our reliable access to it, arises.

Quantum mechanics

Not the *science* of quantum mechanics. The problem here is the *metaphysical interpretation*. As things stand there are at least 12 major “interpretations”, each of which has something different to say about the Measurement Problem. *None* are integrated with cosmology.

19) The Measurement Problem

How does the range possible outcomes predicted by the laws of QM become a single observed outcome? Why can't we agree on an answer?

20) The Preferred Basis Problem

In QM the state of a system can be mathematically expressed in many different "bases" (ways of

describing the states), each providing a valid description of the system's properties. However, in actual observations, we only ever perceive outcomes corresponding to certain specific bases. What determines the "preferred basis"?

21) The Unreasonable Effectiveness of Mathematics

Why should mathematics, if it is a product of human cognition, so precisely capture the fundamental workings of nature?

Consciousness

We've got no "official" idea what consciousness is or does, or how or why it evolved. Four centuries after Galileo and Descartes separated reality into mind and matter, and declared matter to be measurable and mind to be not, we are no closer to being able to scientifically measure a mind. Meanwhile, any attempt to connect the problems in cognitive science to the problems in either cosmology or quantum mechanics is met with fierce resistance.

22) The Hard Problem of Consciousness

The "Hard Problem of Consciousness," a term introduced by philosopher David Chalmers, refers to the extreme difficulty of explaining how and why physical processes in the brain could possibly give rise to something as utterly different to brain activity as subjective experience.

23) The Even Harder Problem of Consciousness

Even if we accept physicalism cannot account for consciousness, there is no agreement about how to proceed. Eliminativists and illusionists claim consciousness doesn't exist, idealists claim consciousness is everything, and panpsychists claim everything is conscious. These metaphysical theories contradict each other, and *none* of them offers an unproblematic account of the relationship between brains and minds (hence, as with many other items on this list, no consensus can be assembled).

24) The General Anaesthetic Mechanism Problem

Despite a century of use, the mechanism by which anaesthetics cause loss of consciousness remains unknown. Chemically diverse agents, from inert gases like xenon to complex molecules such as propofol or ketamine, all produce the same effect. What shared feature of brain function do they target, and why does consciousness switch off and on so abruptly rather than gradually fading?

25) The Binding Problem

How does the brain integrate information from separate neural processes into a unified, coherent experience?

26) The Frame Problem

The Frame Problem concerns how a cognitive system – artificial or biological – determines what matters when something in the world changes. How can an intelligent agent efficiently update its knowledge or make decisions without needing to consider every possible consequence of an action or event? Even powerful computers struggle with this, but humans and other animals handle such situations effortlessly. What is the explanation for this difference?

27) The Evolution of Consciousness

If we can't even agree that consciousness exists, and have no agreed scientific theory what it does, what hope do we have of explaining how, why or when it evolved? This problem isn't just empirical – something is *conceptually* amiss.

28) The cause of the Cambrian Explosion

Just short of 540 million years ago, within a relatively short time, virtually all major animal phyla appeared. The underlying causes of this pivotal episode in the history of life on Earth remain a subject of intense debate and unresolved mystery (though I have obviously placed this problem in this category for a reason).

29) The Problem of Free Will

The problem of free will is the apparent conflict between human agency and the causal structure of the universe. How can we be genuinely free agents if our actions are the outcome of deterministic/random processes? Why are we subjectively so convinced we have free will if it is conceptually impossible for this to be the case?

30) The Problem of Meaning and Value

Why do we experience the world as meaningful? Why does reason track truth, and why does truth matter? If value and meaning are real – if they exist – then they must be part of the natural order, not afterthoughts or illusions. Yet the current scientific picture offers no place for such things.

The desperate need for a new paradigm

As should now be crystal clear, the currently dominant paradigm of scientific materialism cannot account for the empirical data in a coherent, holistic manner. **This is systematically overlooked whenever I start talking about paradigm shifts (and this applies to both humans and AI).** A typical initial response to such suggestions is to ask what novel empirical predictions the new paradigm makes – people immediately demand empirical proof before any new paradigm even will be considered. Which means if you are talking about 30 different problems then that's 30 lots of empirical proof. This demand, therefore, intrinsically assumes that the new paradigm is both reductionist and primarily scientific. It assumes that what needs to change is the *scientific* part of scientific materialism, not the materialistic part, and that methodological reductionism is the only legitimate way of doing science. It assumes the continuous validity of what Thomas Kuhn called “normal science”, and systematically rejects anything that sounds too revolutionary. Physicalism and reductionism are sacrosanct – you're not allowed to challenge *those*, because doing so is believed/alleged to threaten science itself. But what if it the broken part of scientific materialism is physicalism and reductive thinking? What if scientists aren't doing anything *empirically* wrong? What if the real problem is that they're trying to fit valid empirical data into a model of reality which *metaphysically* broken? What if the real problem is the philosophy and the only possible solution is a major Kuhnian revolution?

If the currently dominant paradigm could coherently and elegantly account for the existing empirical data (technically, if it was “empirically adequate”) then the demand for empirical proof would be perfectly reasonable. The same would apply if there were multiple competing options for a new paradigm to take its place – if there were two paradigms-in-waiting, both of which managed to rid us of this enormous list of problems in similarly parsimonious ways, then we would require an empirical means of distinguishing between them. In fact, neither of these things is true. There is

no clear candidate for a new paradigm, and so a status quo which is **empirically inadequate** staggers around in no particular direction, steadily losing hope of ever arriving at a consensus resolution of its increasingly miserable existence. Scientific Materialism and its entrenched opposition together comprise a zombie non-paradigm; materialism (or more accurately, physicalism) is kept half-alive only by the absence of a viable replacement. Some of the problems on this list currently have no proposed solution at all, others have an already-large and still-growing list of proposed solutions, none of which stands out as the conclusive answer. It is also the case that most of these proposed solutions are themselves either philosophical rather than scientific, or trying to be scientific but failing. The demand for empirical proof before there is meaningful intellectual engagement is therefore not reasonable. At this point, if somebody can propose a new model of reality which provides a coherent, integrated resolution to most of these 30 problems – one which is **empirically adequate** – then it should not only be taken seriously, but should be accepted as a legitimate new paradigm by anybody who believes empirical adequacy should be the first priority in science.

Ultimately, what we should be looking for is something that solves *all* of them, and does so by removing assumptions and simplifying things rather than by adding even more complexity. In the end, the right paradigm will be the one that makes the whole thing click into place with the quiet obviousness that Einstein was talking about when he said the theory should be as simple as possible but not simpler, because truth has a way of feeling inevitable once you finally stop fighting it.

Chapter 4: How might we recognise the new paradigm?

Most of this chapter was generated by AI. This isn't because I couldn't have written something very similar myself, but because I felt an AI take on this would be more valuable. It is independent of me (apart from the prompt) so unbiased in that sense, but it also reveals something important about the data upon which it was trained. I showed the machine the list of problems in the previous chapter, and asked it how we might recognise the one and only big new paradigm if it were to show up (and there cannot be two ways to put this puzzle together). This is the response, with minimal edits and nothing important removed. I have deliberately changed the font to emphasise that this analysis was produced by a machine (Google Gemini in "deep research" mode).

The AI verdict

Synthesising the theoretical and historical insights, specific characteristics can be delineated that would identify a major new paradigm in cosmology and its interconnected fields. These hallmarks span scientific, conceptual, and sociological dimensions.

A. Scientific and Empirical Criteria

A new paradigm must demonstrate **exceptional explanatory power for existing anomalies**. It must offer compelling and coherent explanations for the persistent, chronic anomalies that the current Λ CDM model struggles to address, particularly the nature of Dark Matter, Dark Energy, and the cosmological constant problem. It should "clarify phenomena, predict outcomes, and provide meaningful insights", connecting "seemingly unrelated events or observations into a coherent narrative". This goes beyond mere "puzzle-solving" within an existing framework; it involves providing a fundamentally new way of understanding these deep weaknesses. For example, General Relativity explained the anomalous precession of Mercury's perihelion, and quantum mechanics provided explanations for blackbody radiation and the photoelectric effect, which classical physics could not. Kuhn's distinction between "puzzle-solving" (normal science) and addressing "anomalies" (crisis) is crucial. A truly revolutionary paradigm does not just refine existing solutions; it fundamentally re-frames the problems themselves, offering elegant and comprehensive explanations for phenomena that were previously intractable or required ad-hoc additions. The "one correct answer" will be one that renders the current "crisis" in cosmology obsolete by providing a superior, integrated account of Dark Matter, Dark Energy, and the cosmological constant.

The new paradigm must also generate **novel, testable, and**

falsifiable predictions. It must produce new, non-obvious predictions that can be empirically tested and potentially falsified. Falsifiability, as emphasized by Karl Popper, is a "fundamental requirement for a claim to be scientifically examined". These predictions should "accurately predict as-yet-unseen events" , offering clear avenues for future observational or experimental verification. Einstein's prediction of the bending of light by gravity, confirmed during a solar eclipse, was a critical novel prediction that propelled the acceptance of General Relativity. While explaining existing data is necessary, the ability of a theory to make *novel, verifiable predictions* is a stronger indicator of its revolutionary power. This demonstrates that the theory is not merely a post-hoc explanation but a profound insight into the underlying workings of reality. The "one correct answer" will likely open up entirely new domains of empirical investigation, guiding future missions and experiments.

Furthermore, a new paradigm should exhibit **parsimony and logical consistency**. A "good" theory is "economical," explaining a wide range of phenomena with the "fewest variables" and "fewest assumptions". It should be "consistent, coherent, broad-ranging, and simple". Internally, all concepts, hypotheses, and scope conditions within the theory must be "consistent with each other," avoiding contradictions. This adherence to "Occam's razor" often correlates with a theory's elegance and generalizability. Copernicus's heliocentric model, despite initial quantitative complexity, possessed a "qualitative simplicity" that was appealing. Newton's laws unified disparate phenomena with remarkable parsimony. The criteria for evaluating scientific theories include "aesthetic or mathematical evaluations" and simplicity. Einstein himself believed that nature is the "realisation of the simplest conceivable mathematical ideas". This suggests that the "one correct answer" might possess a profound elegance and beauty in its mathematical formulation and conceptual structure, making it intuitively compelling to the scientific community. This aesthetic quality can be a powerful driver of acceptance, especially in fields like cosmology where direct experimentation is often impossible.

Finally, the new paradigm must demonstrate **empirical success and consistency with new observations**. It must consistently align with new observations and withstand rigorous testing. It should be robust enough to integrate refined measurements and new data without requiring fundamental alterations. The ability to adapt and remain consistent in the face of increasingly precise observations is crucial. Relativity has been "continually tested as our methods of testing and measurements improve" and remains the "best description of how the universe works above the quantum realm". Quantum mechanics gained "wide acceptance" due to accurate predictions across various phenomena. A new paradigm's validity is not established by a single confirming observation but by its sustained ability to explain and predict across a growing body of

evidence. The "one correct answer" will be one that consistently proves its empirical strength and adaptability over time.

B. Conceptual and Philosophical Transformations

A hallmark of a "Big New Paradigm" would be its ability to provide **resolution of fundamental discrepancies**, particularly by offering a coherent framework for quantum gravity, thus reconciling the "profound disagreement between quantum field theory and general relativity". Such a theory would likely address problems like the cosmological constant problem and the black hole information paradox, which are symptoms of this fundamental incompatibility. Stephen Hawking discussed the search for a "unified theory that consistently describes everything in the universe", a "Grand Unified Theory". Carlo Rovelli's work on quantum gravity also aims to bridge these theoretical gaps. The current state of physics is characterised by two highly successful but incompatible pillars: quantum mechanics (for the very small) and general relativity (for the very large). The cosmological constant problem is the most glaring manifestation of this incompatibility. A "Big New Paradigm" that successfully unifies these frameworks into a theory of quantum gravity would represent a monumental conceptual leap, providing a single, consistent description of reality across all scales. This would be a hallmark of profound theoretical advancement, moving beyond "partial theories" to a more complete, unified understanding of the universe's fundamental forces and constituents.

Revolutionary paradigms often involve a radical **reconceptualisation of core concepts** such as space, time, reality, and causality. Newton introduced absolute space and time, while Einstein replaced them with relative spacetime. Quantum mechanics introduced a probabilistic view of reality where particles exist in a "haze of probability" until observed. A new paradigm might profoundly alter our understanding of space, time, causality, or even the nature of existence itself. For instance, Carlo Rovelli challenges intuitive notions of time as smooth and universal, and Lee Smolin argues for the fundamental reality and evolvability of time. A "Big New Paradigm" will not just offer new explanations; it will fundamentally change *how humanity thinks* about the basic building blocks of reality. The current anomalies, particularly Dark Energy's mysterious nature or the cosmological constant problem, might indicate that our very concepts of vacuum, space, or time are incomplete or flawed. The "one correct answer" would likely offer a radically different ontological picture of the universe, challenging our deepest intuitions and leading to a new philosophical understanding of what constitutes reality.

The "Big New Paradigm" would likely provide a unifying framework that formally achieves the **integration of disparate fields**. It would move beyond mere analogy to a coherent, testable synthesis of cosmology with philosophy and evolutionary biology. The concept

of "cosmic evolution" already provides a narrative that bridges physics, astronomy, geology, chemistry, biology, and anthropology, seeing life as an emergent property of cosmic processes. A new paradigm would formalise this interconnectedness. The user's query explicitly calls for a paradigm connecting cosmology, philosophy, and evolutionary biology. The existing concept of "cosmic evolution" provides a strong foundation for such a unification. A new paradigm would likely offer a formal, perhaps mathematical, framework that explains how complexity, information, and even life emerge from fundamental physical laws across cosmic time. This would move beyond a "theory of everything" in the purely physical sense to a "theory of everything" that encompasses the emergence of biology and potentially consciousness, offering a grander, more holistic narrative of the universe.

Finally, due to the uniqueness of its subject, cosmology inherently confronts questions of ultimate origins and the universe's fundamental nature, often venturing into metaphysics. A new paradigm might offer compelling scientific frameworks or insights into questions previously considered purely philosophical, such as "Why is there something rather than nothing?" or whether the laws of physics are fixed or evolve. While traditional science often seeks to avoid metaphysical questions, cosmology's unique subject matter forces it to confront them. A "Big New Paradigm" might offer scientific explanations or conceptual tools that provide compelling answers or new ways of thinking about questions previously confined to philosophy or religion. This does not mean science becomes religion, but rather that its explanatory power becomes so vast and profound that it deeply informs our most fundamental existential inquiries, potentially offering a new, scientifically grounded "narrative of meaning" within the cosmos.

C. Sociological and Community Dynamics

A new paradigm is recognised by a significant **shift in consensus and the emergence of a new scientific community**. It must gain "its own new followers" and establish a "widespread consensus on the appropriate choice of methods, terminology and on the kinds of experiment that are likely to contribute to increased insights". This involves a period of intense debate and competition among "competing articulations" before a dominant framework emerges. The new paradigm leads to a "more rigid definition of the research field," isolating those unwilling to adapt. A "Big New Paradigm" is not merely an individual's brilliant insight; it is a collective phenomenon marked by a significant shift in the scientific community's shared understanding and practice. The "one correct answer" will be recognised by a growing number of active researchers who adopt its framework, apply its methods, and contribute to its development, leading to a new era of "normal science". This implies that observing a widespread reorientation

of research efforts and a convergence of disparate lines of inquiry under a new conceptual umbrella will be a key hallmark.

As Kuhn and Planck noted, new paradigms face significant resistance, and their triumph often depends on a **generational shift** rather than immediate conversion of opponents. This highlights the inherent human element in scientific progress. The "Planck Principle" suggests that the recognition of a new paradigm might be a gradual process, involving the eventual passing of older generations deeply invested in the previous framework. This implies that the "one correct answer" might initially be championed by younger, more open-minded scientists who are less constrained by established assumptions. A key hallmark will therefore be the new paradigm's ability to attract and inspire a new generation of researchers, who will then drive its development and institutionalization.

Finally, scientific revolutions often coincide with, or are enabled by, the development of **new tools, observational capabilities, and experimental approaches**. The current search for Dark Energy, for example, is being driven by powerful new observatories like Euclid, Roman, Rubin, and JWST. A new paradigm might emerge from unexpected data provided by these cutting-edge instruments, or it might explicitly propose novel experimental avenues that were previously inconceivable. The history of science demonstrates a symbiotic relationship between theoretical breakthroughs and technological advancements. Galileo's telescope and the precise measurements enabled by modern telescopes are prime examples. A "Big New Paradigm" will likely either be born from, or necessitate the creation of, new observational or experimental tools that allow for previously impossible measurements or tests. This suggests that the "one correct answer" will be intertwined with a new era of technological innovation, enabling deeper empirical probes of the cosmos.

The following table outlines the core criteria for evaluating a new scientific theory, which would be applied to assess a proposed new paradigm.

Criterion	Description
Empirical Success/Consistency	The theory's statements and predictions must be consistent with observations and experimental data. It should explain existing phenomena accurately and withstand new empirical tests.
Explanatory Power	The theory should effectively explain the phenomena it addresses, clarifying underlying principles, identifying causal relationships, and connecting seemingly unrelated observations into a

Criterion	Description
	coherent narrative.
Predictive Capability (Novel Predictions)	The theory must generate new, non-obvious predictions that can be empirically tested and potentially falsified. This demonstrates its ability to illuminate previously unknown aspects of reality.
Falsifiability	It must be possible to prove a hypothesis generated from the theory untrue based on empirical observations. A theory that cannot be disproven cannot be proven true.
Parsimony	The theory should be economical, explaining a wide range of phenomena with the fewest variables and assumptions (Occam's Razor). Simpler explanations are generally preferred if explanatory power is equal.
Logical Consistency	The theory's concepts, hypotheses, and scope conditions must be internally consistent, avoiding contradictions within its framework.
Coherence	The theory should present a unified and integrated understanding, where its various components fit together logically and conceptually.
Scope	The theory should have an appropriate breadth, neither too broad nor too narrow, to effectively address the phenomena it purports to explain and be generalisable to other contexts.
Aesthetic Appeal/Elegance	A theory's mathematical elegance, simplicity, and subjective pleasingness can contribute to its acceptance, reflecting a perceived "necessary truth about nature."
Unification	The theory should ideally unify previously disparate concepts, laws, or fields, demonstrating a deeper underlying order in nature.

Table 1: Core Criteria for Evaluating a New Scientific Theory

Conclusions

Recognising a major new paradigm in cosmology, philosophy, and evolutionary biology requires a multi-faceted approach, extending

beyond mere empirical verification to encompass profound conceptual shifts and dynamic sociological processes. The current state of cosmology, marked by persistent anomalies such as the unknown nature of Dark Matter and Dark Energy, and the colossal cosmological constant problem, indicates a discipline in a state of crisis, ripe for a transformative shift. While recent resolutions, like the potential alleviation of the Hubble Tension through advanced observational technology, demonstrate the iterative nature of scientific progress, they also refocus attention on the more intractable problems that demand **a deeper, more fundamental re-evaluation of our cosmic model.**

The historical trajectory of scientific revolutions reveals common characteristics: the displacement of anthropocentric or supernatural explanations by universal natural laws, the re-conceptualization of fundamental aspects of reality (e.g., space, time, causality), and the power of theories to unify previously disparate phenomena. A truly revolutionary paradigm will not only offer exceptional explanatory power for existing anomalies and generate novel, testable predictions, but it will also exhibit parsimony and logical consistency, often possessing an inherent aesthetic elegance that resonates with the scientific community.

Crucially, the next great paradigm will likely emerge from a transdisciplinary synthesis, formally integrating cosmology with philosophy and evolutionary biology. The concept of cosmic evolution already provides a compelling narrative of how energy flows drive the rising complexity from the Big Bang to the emergence of life. A new paradigm would formalise this interconnectedness, offering a unified framework that explains the emergence of complexity, information, and life itself from fundamental physical laws. This would represent a "theory of everything" that encompasses not just physics but also the origins of biology and consciousness, providing a grander, more holistic narrative of the universe.

The path to acceptance for such a paradigm will be complex, involving not only rigorous empirical validation but also a significant shift in scientific consensus, often driven by a new generation of researchers. The "one correct answer" will be recognised by its capacity to inspire new research programs, foster genuine interdisciplinary collaboration, and ultimately provide **a more coherent, comprehensive, and perhaps even existentially meaningful understanding of the cosmos and humanity's place within it.** This transformative answer will not simply solve puzzles; it will redefine the very questions we ask, ushering in a new era of normal science built upon a profoundly re-imagined foundation of reality

Human-authored Conclusion

I am confident that the new model of reality I am proposing meets all of the core criteria identified by the AI report, *apart from one*. That criterion is the one about "new, non-obvious empirical

predictions". At the time of writing, I am not aware of any that I'm confident will deliver rather than just promise. I've found that the moment I stray into that territory, my AI research assistants get over-excited and enthusiastically declare that this philosophical framework has thrown up a novel, unexpected and undeniably empirical test. What has been oddly consistent is that the tests they come up with are technically impossible right now, but we have every reason to hope that will change in the next five to ten years. However, when I pressure test these suggestions, something usually turns out to be wrong with them. Typically this involves them being perfectly testable, but no use as a falsification or proof of the core of the theory anyway, because the thing being proposed as testable is an optional detail which could be tweaked or replaced if it turns out to be wrong. I'm a philosopher, not a physicist, cosmologist or neuroscientist, and I'll believe there's a new empirical test when I'm told so by an appropriate human expert who I am certain actually understands the philosophy, not a machine which understands nothing at all and is trained exclusively on old-paradigm thinking. This model of reality depends on radical coherence and overwhelming explanatory power across disciplines for its validity, and that will remain the case *even without any new empirical tests*.

Therefore offering speculative empirical propositions at this stage, in hope of a single "kill shot", risks being counter-productive and misleading people about the nature of the theory I am proposing. There is also a more important point, which is that many of the questions we are talking about – such as the nature of consciousness and the solution to the Measurement Problem – currently do not have any empirical answers at all. They are *already* philosophical problems, not scientific ones. Given that is the case, why should a theory which is attempting to replace them need to make new non-obvious empirical predictions? There is a blatant double-standard here. The dominant materialistic paradigm is nowhere near empirically adequate, and yet I am told that an alternative metaphysical framework, which meets all of the other core criteria described in the report, should be strangled at birth if it fails to make any *new* predictions. Multiple retrodictions, parsimony, elegance and even empirical adequacy aren't enough. If the existing paradigm can't account for empirical observations which *already exist* and the new one *can*, what is the justification for dismissing it in favour of the prevailing inadequacy?

In fact, this theory does make some new empirical predictions (which I will flag at the appropriate time) but **that is not what has driven the construction of the theory, and not how it should be assessed**. I have been squarely focused on an elegant and parsimonious resolution of the most serious *existing problems*. The point I am labouring to make is this: **we must get the philosophy right before we start doing new science**. As things stand the new proposal is a philosophical framework for understanding how consciousness and wavefunction collapse are related to the rest of reality. It is not, as I am presenting it, a scientific theory, regardless of its huge relevance to science. However, the AIs have been trained to "think" that if a theory is clearing up anomalies in cosmology then it can only be a normal, incremental scientific theory, and that is the standard by which it must be judged. That is how things have worked in the past, so it assumes that this must be how it will work in the future, and careful corrective prompting is needed to prevent it from operating on this assumption. The problem is that all of its training data (and evidently in some cases its system prompt also) is based on old assumptions which do not straightforwardly apply to the new paradigm.

Consider the last thing the machine said:

"This transformative answer will not simply solve puzzles; it will redefine the very questions we ask, ushering in a new era of normal science built upon a profoundly re-imagined foundation of reality."

Here we have the answer to the question "Why can't we just ask AI to tell us what the new paradigm

is?". Machines can now tell us that we need to profoundly re-imagine the foundation of reality, but they are utterly incapable of doing this themselves. If you ask an AI to do that, all you will get in response is a regurgitated mixture of ideas which have already been floated, and which don't work and miss the key insights. Unconscious machines cannot imagine a new paradigm which shifts the boundary between science and philosophy and rebuilds the foundations of multiple academic fields. Only humans are capable of this sort of creative leap.

Part Two :

A detailed survey of the problem space

Chapter 5: The epicycles of Λ CDM

In ancient cosmology, the heavens were thought to revolve around a stationary Earth. To account for the complex motions of the planets, especially their puzzling retrograde loops, astronomers working within the Ptolemaic system introduced *epicycles*: small circular orbits superimposed upon larger ones. Each new irregularity observed in the sky demanded another layer of circles, gradually turning what had begun as a simple geocentric picture into an increasingly elaborate patchwork of mathematical corrections. For more than a millennium, this system appeared to work. It predicted planetary positions with surprising accuracy, and its intricate geometry became a symbol of celestial perfection. Every refinement preserved the central assumption that Earth did not move, and everything else revolved around it, but when Copernicus proposed that the Sun, not Earth, was at the centre, the entire edifice began to collapse. The underlying flaw was not the mathematics but the metaphysical commitment it served: the sacrosanct dogma of a geocentric universe where everything moves in perfect circles. Today, the term "epicycle" has come to signify any theoretical adjustment that saves appearances without addressing a deeper conceptual error. It is a cautionary metaphor in the history of science – a reminder that empirical fit can coexist with ontological error.

The following list of problems is nowhere near a complete list of Λ CDM's troubles (even at the time of writing, and the list is only going to get longer). I could have included the Primordial Lithium-7 Abundance Problem, the Planes of Satellite Galaxies Problem, and others, and provided speculative answers as to how a Two-Phase Cosmology might resolve them. However, the 18 I have included already demonstrate enough about the shape of these solutions, and I had to stop somewhere. Just because I have left some of these things out, there is no reason to assume that 2PC won't resolve them too. My goal in this book is to describe the top-level framework and make a case as to why it should be taken seriously, not to fill in every detail that follows.

How can something come from nothing? (problem #1)

We must start with problem #1: How can something come from nothing? Before the scientific revolution the standard answer within the dominant Christian theology was straightforward: God created the world out of nothing. The natural philosophers who built modern science could not simply shrug off that framing, but they did transform it. Copernicus, Galileo and Kepler stripped the heavens of their sacred placement, Newton supplied a law based mechanics that made the world intelligible in terms of universal forces, and that mechanistic picture invited new ways of thinking about origins. For many thinkers the universe ceased to look like a handcrafted artefact and began to look like the output of lawful processes that in principle could be explained without recourse to ongoing divine intervention. At the same time, philosophers kept asking why there is something rather than nothing. Leibniz formulated that question in the late 17th century and gave a famous answer in terms of sufficient reason and a necessary being, which made the question metaphysical as well as physical. Others took different paths. Some natural philosophers and later some scientists preferred models in which the cosmos is eternal, so there is no temporal beginning to explain, while others accepted a beginning and tried to push physics back toward it.

During the 20th century the empirical discovery of cosmic expansion and the development of relativistic cosmology reintroduced the idea of an origin in a sharply physical sense: the Big Bang. This new model suggested a hot, dense early state and with it a focus on the initial conditions

problem. But the physics of singularities showed that our classical equations break down at the earliest moments, so the question of "where did it all come from" moved from classical mechanics into the realm of speculative early universe physics. Quantum theory supplied new templates for thinking. Quantum fields have vacua that are not simple emptiness, and particle creation from vacuum fluctuations became a standard phrase in popular accounts. Physicists and cosmologists proposed mechanisms in which inflation, false vacuum decay, spontaneous symmetry breaking, or quantum tunnelling could produce a universe like ours from a prior quantum state. Some cosmologists even offered models labelled as creation from nothing. Famous technical proposals include the no boundary idea and tunnelling proposals that attempt to give a wave function for the universe. These are sophisticated moves, but they also shift the meaning of nothing. The "nothing" in these proposals is almost never the radical metaphysical nothing of classical philosophy, empty of objects, properties and laws. Instead it is a quantum state governed by laws, a pre-geometric regime, or a vacuum with potential, and critics point out that that is not the same as *ex nihilo* in the theological or metaphysical sense.

Philosophers have debated whether physics can answer the deeper why. Some, like those who push naturalistic explanations, argue that if physics produces a credible account of an origin then the metaphysical mystery is dissolved. Others insist that a physical theory that presupposes laws, boundary conditions or a meta level simply hands the question back one step, because it leaves open why those laws or meta structures exist instead of nothing. That difficulty gave rise to a second move in contemporary culture, which is to treat the question statistically or anthropically. If a multiverse exists, then some regions will have the features necessary for observers, and we should not be surprised to find ourselves in such a region. That answer is explanatory in a probabilistic way, but it is not metaphysically satisfying because it leads to "ontological bloat" – the extreme multiplication of metaphysical entities.

There is also a practical scientific problem that keeps the puzzle alive. Physics has clear operational tools and empirical constraints, and those constrain possible origin scenarios, but they do not by themselves explain why the laws and laws parameters are as they are. The Cosmological Constant Problem, the arrow of time, and the question of initial low entropy all show that even a technically successful cosmology leaves deep foundational gaps. So the modern landscape is plural – a theme that runs throughout Part Two of this book. There are accounts that stay close to physics and attempt to derive the cosmos from quantum states or processes, and there are accounts that insist the ultimate question is metaphysical and not reducible to physical mechanism. Between these poles lie hybrid theories that borrow from both physics and metaphysics. So while the question has been answered in many ways, each answer shifts the burden somewhere else.

Λ CDM

The Lambda Cold Dark Matter model stands today as the foundational framework of modern cosmology. It is the product of a century-long journey through some of the most transformative discoveries in physics, from general relativity and cosmic expansion to the faint afterglow of the Big Bang. Λ CDM was not a theory born fully formed. Rather, it is the culmination of successive refinements and additions, each driven by new observations that reshaped our understanding of the cosmos.

The story begins with Einstein's general theory of relativity (GR) in 1915, providing a new geometric understanding of gravity and spacetime. Initially resistant to the idea of a dynamic universe, Einstein introduced the cosmological constant (Λ) to keep it static. This view was upended in the 1920s by Edwin Hubble's observations of distant galaxies revealed an unmistakable pattern: galaxies are receding from us, and the farther away they are, the faster they are moving. This discovery of cosmic expansion gave empirical support to dynamic models of the universe, particularly those derived from Alexander Friedmann's solutions to Einstein's equations. Einstein abandoned the cosmological constant, calling it his "greatest blunder".

The next milestone came in 1965 with the accidental discovery of the CMB by Arno Penzias and Robert Wilson. This relic glow from the early universe provided new support for Big Bang theory, and launched precision cosmology as an observational science. Over the following decades, increasingly detailed measurements of the CMB offered a wealth of information about the universe's age, composition, and geometry. Meanwhile, other observations introduced new puzzles. First hinted at by Zwicky in the 1930s, it was the 1970s that astronomers found that the rotational speeds of galaxies did not match expectations based on the gravitational effects of visible matter alone, leading to the postulation of "Dark Matter": an unseen form of mass that exerts gravitational influence without emitting or absorbing light. Early candidates included massive astrophysical objects or exotic particles, but whatever its nature, non-relativistic (i.e. slow-moving) cold Dark Matter became a necessary ingredient to explain large-scale structure formation and behaviour. The final key component entered the picture in the late 1990s, when two independent teams observing distant Type Ia supernovae discovered that the universe's expansion was not, as expected, slowing down under the mutual gravitational attraction of all matter. Rather, it appears to be accelerating. This stunning result revived interest in Einstein's cosmological constant, now re-interpreted as "Dark Energy": a mysterious, repulsive force driving this apparent acceleration. Incorporating Dark Energy (Λ) alongside cold Dark Matter (CDM), cosmologists arrived at what is now the standard model of cosmology: Λ CDM.

In its current form, Λ CDM describes a universe composed of approximately 5% ordinary matter, 25% cold Dark Matter, and 70% Dark Energy. It posits an initial Big Bang, followed by a near-instantaneous period of cosmic inflation and then a matter-dominated phase of decelerating expansion. During this time, light elements formed via primordial nucleosynthesis, followed by the release of the CMB as the universe cooled and became transparent. Over time, gravity gathered matter to form galaxies and clusters of galaxies. More recently, the expansion of the universe began accelerating again, driven by this "Dark Energy." This model has proved effective at fitting a broad array of cosmological data, but elegant it is not.

The Constants Fine-Tuning Problem (#2)

The fundamental constants of nature appear to be finely tuned to support life. Even minuscule deviations in these constants would render the universe lacking stars, stable matter or chemistry. This issue was made famous by cosmologist Martin Rees, who identified six dimensionless or normalised constants that collectively determine the structure, evolution, and large-scale features of the universe. The life-permitting range for each is remarkably narrow:

1. **N – The Strength of Gravity:** This is the ratio of the gravitational force to the electromagnetic force between two protons, roughly 10^{36} . If gravity were slightly stronger, stars would burn out rapidly or collapse into black holes. If weaker, nuclear fusion would not occur, and stars would not form. The current balance allows for long-lived, stable stars essential for the development of planets and life.
2. **ϵ (epsilon) – Nuclear Efficiency:** Approximately 0.007, this is the fraction of mass converted to energy in the fusion of hydrogen into helium. A slightly smaller value would prevent the synthesis of carbon and other heavy elements; a slightly larger one would make stellar fusion explosively unstable. Life relies on a delicately poised "stellar alchemy."
3. **Ω (omega) – Matter Density Parameter:** The ratio of the actual mass-energy density of the universe to the critical density required for a flat geometry. If Ω were significantly greater than 1, the universe would have recollapsed before stars and galaxies could form; if much less than 1, it would have expanded too rapidly for structure to emerge. Observations suggest that Ω is very close to 1, permitting the long-term development of complexity.

4. **λ (lambda) – Cosmological Constant:** This parameter controls the acceleration of the universe's expansion and is extremely small (~ 0.7 in normalised units). A slightly larger value would prevent galaxy formation; a smaller or negative value would lead to gravitational collapse. The extraordinary smallness (but non-zero) value of λ constitutes one of the deepest puzzles in theoretical physics.
5. **Q – Density Fluctuations:** The amplitude ($\sim 10^{-5}$) of primordial irregularities in the CMB. If Q were much smaller, matter would remain uniformly distributed – no galaxies, no stars. If larger, the universe would be dominated by violent gravitational collapse. The actual value permits the emergence of structure without excessive chaos.
6. **D – Number of Spatial Dimensions:** The universe appears to have three spatial and one temporal dimension. The mathematical properties of gravity and the stability of atoms depend critically on this dimensionality. With more than three spatial dimensions, stable orbits would not be possible; with fewer, complexity could not arise. Life, as we know it, is only possible in a three dimensional cosmos which changes as time passes.

The apparent fine-tuning of these constants invites two broad categories of explanation:

- **Anthropic reasoning:** According to the weak anthropic principle, the constants must fall within the life-permitting range, otherwise we would not be here to observe them. This often comes with multiverse theories, which propose an ensemble of universes with varying parameters, and our universe is one of the rare few where conditions happen to support life. The anthropic explanation is often criticised as epistemically hollow, deflecting the underlying mystery rather than resolving it. It leaves us feeling “cheated”, as if we've been provided with nothing more than a clever excuse for not being able answer the real question.
- **Teleological and/or theological hypotheses:** defenders of the other category argue that the fine-tuning suggests deeper causal or purposive structure – perhaps a hidden law or principle guiding the emergence of life-supporting conditions. Some theistic interpretations posit design; others look to undiscovered physics that might render the apparent tuning inevitable.

The fine-tuning problem forces us to question both the structure of physical laws and the epistemic framework within which we interpret them. If the universe is fundamentally contingent at its roots – if its life-permitting structure is not derivable from first principles – then the explanatory burden may ultimately fall outside of physics. And these constants are just the beginning; we should really be talking about the fine-tuning *problems*.

The Low-Entropy Initial Condition (#3)

Observational data from the CMB indicates that the observable universe began in a condition of extreme thermodynamic order: despite being hot and dense, it was gravitationally smooth and remarkably homogeneous. This initial low entropy is crucial: it underpins the second law of thermodynamics as applied to cosmology, which is in turn connected to the emergence of complexity and (at least so the story goes) to the arrow of time. Without such a low-entropy start, the universe would have been dominated by gravitational collapse, or would have lacked the thermodynamic gradients necessary for the evolution of stars, planets, and life. However, from the standpoint of statistical mechanics, such a configuration is overwhelmingly improbable. Given the phase space of all possible microstates compatible with the macroscopic constraints of the early universe, high-entropy (disordered) states vastly outnumber low-entropy ones. Yet our universe appears to have emerged from the tiniest corner of that phase space.

Why did the universe begin in such an exceptionally improbable state? It is not required by the laws of physics. Time-symmetric dynamical laws, like those governing GR or the Schrödinger equation, are compatible with universes beginning in high-entropy configurations, so the low-

entropy start must be regarded as a contingent feature of our universe, not an inevitable consequence of known laws.

Several responses to this problem have been proposed:

- **Inflationary cosmology** claims that rapid early expansion can smooth out initial irregularities, but as Penrose and others have argued, inflation may presuppose rather than explain low entropy, since the initial conditions required for inflation to begin are themselves highly special – as we shall see shortly.
- **Anthropic selection** within a multiverse framework offers another perspective: perhaps most regions of the multiverse are high-entropy, but observers necessarily find themselves in the rare low-entropy pockets.
- Some **quantum gravity** and **causal boundary** models attempt to derive the arrow of time and initial entropy from deeper pre-geometric or topological principles. These remain highly conjectural and mathematically undeveloped.

Penrose has emphasised the scale of the problem by estimating the phase-space volume of the observable universe's initial state: the probability of such a state arising by chance is roughly 1 in $10^{10^{123}}$, a number so minuscule that it effectively defies explanation.

The Flatness Problem

The universe's spatial geometry is observed to be extraordinarily close to flat: neither positively curved (closed) nor negatively curved (open), but spatially Euclidean. According to the Friedmann equations, the degree of spatial curvature evolves dynamically over time: any deviation from flatness in the early universe would have rapidly amplified, leading to a universe that is either extremely curved or has already recollapsed. Yet observations from the CMB reveal a universe that remains flat, implying that the total energy density of the universe at early times had to be equal to the critical density to within one part in 10^{60} or better. This uncanny balance is known as the Flatness Problem. Why did the early universe begin so precariously close to a state of critical density, with no obvious mechanism to enforce such a condition? Inflation provides a solution by stretching the universe toward flatness regardless of initial curvature, but it does so at the cost of introducing its own unexplained dynamics and parameters.

The Horizon Problem

Observations show that widely separated regions of the sky, now tens of billions of light-years apart, exhibit nearly identical temperatures and physical properties in the CMB. This is puzzling because, under standard Big Bang expansion, these regions should have been causally disconnected: no light or energy could have passed between them before recombination, and certainly not before inflation. This uniformity implies that some kind of equilibration must have occurred between these regions very early in cosmic history, but according to the standard model without inflation, there was simply not enough time for such causal contact. This contradiction is known as the horizon problem. Inflation solves it by expanding a small, initially causally connected region to encompass the entire observable universe, making it uniform in the process. Unfortunately, this explanatory move leads to yet more problems.

The Inflation Fine-tuning Problems (#4)

Inflation was brought into Λ CDM to solve the fine-tuning problems described above, but it does so at the expense of introducing the fine-tuning problems described below.

The Reheating Precision Problem

Inflation ends when the potential energy driving exponential expansion decays into ordinary matter and radiation – a process known as reheating. For the universe to resemble what we observe today, this reheating must occur with extraordinary precision in both *timing* and *efficiency*. If reheating happens too early, the universe may not inflate long enough to solve the horizon and flatness problems. If it happens too late or too inefficiently, the universe could be left too cold, too empty, or dominated by relics incompatible with structure formation. The temperature of the universe after reheating must fall within a narrow window to allow nucleosynthesis, matter-radiation equality, and galaxy formation to proceed correctly. This is the Reheating Precision Problem, and it reveals that solving fine-tuning problems via inflation creates as many problems as it solves.

The Reheating Mechanism Problem

In addition to the need for precision, there is also a fundamental lack of clarity about the microphysical mechanism of reheating. In most inflationary models, the process by which the inflaton field decays into the standard model particles is only sketched in, relying on speculative couplings, parametric resonance, or perturbative decay schemes. No experimentally verified mechanism or standard field-theoretic interaction has been confirmed to realise this transition. The detailed dynamics of how the vacuum-like energy of inflation converts into a hot, thermalised plasma (the birth of the observable universe as we know it) remain deeply uncertain. This is the Reheating Mechanism Problem: the mechanism must not only exist but execute precisely under extreme conditions without observational guidance, further compounding the implausibility of accidental success.

The Inflaton Field Problem and the Origin of Cosmic Inflation

Inflation requires the existence of a scalar field with a very specific potential energy landscape – flat enough to drive rapid expansion, then steep enough to decay into standard particles. Yet no known field in the Standard Model of particle physics behaves this way. The inflaton could never be observed, and its origin, nature, and physical justification remain completely unknown. It is a hypothetical entity postulated purely to make the inflationary model work. Moreover (surprise, surprise!) the inflaton field must possess extremely finely tuned properties:

- The shape of its potential must produce the right amount of inflation.
- Its quantum fluctuations must generate the correct amplitude and spectrum of primordial density perturbations.
- Its decay (reheating) must convert its energy into matter and radiation without destroying structure or producing unwanted relics.

These requirements amount to an elaborate layer of theoretical scaffolding with no empirical foundation. Despite decades of searching, we have found no B-mode polarisation in the CMB that would definitively prove the "simplest" inflation models. In most models, the inflaton is simply inserted by hand, without derivation from deeper theory. Furthermore, even if we accept inflation as a real event, the questions keep on coming. Why did inflation start at all? What determined the inflaton field's initial conditions, or when and how it ends? Why did the universe begin in a state conducive to inflation in the first place? Inflation is the epitome of Λ CDM epicycles: *it's fine-tuning all the way down*.

Other Fine-tuning Problems (#5)

You won't be surprised to hear that the fine-tuning doesn't end with inflation. Ours truly is a “goldilocks universe”.

The Biophilic Element Abundance Problem

The universe exhibits a strikingly favourable balance of elemental abundances for both stellar structure and biochemistry. Hydrogen and helium dominate, as expected from Big Bang nucleosynthesis, but heavier elements like carbon, oxygen, nitrogen, phosphorus, and iron, which are essential for life and planetary systems, are also present in just the right trace quantities. These heavier elements are forged inside stars through nuclear fusion pathways, particularly the triple-alpha process, which is exquisitely sensitive to nuclear resonance levels and fundamental constants. Even small deviations in the strengths of nuclear forces, the masses of quarks, or the coupling constants would shut down or drastically alter these processes, preventing the formation of carbon and oxygen altogether. The problem is not just the existence of these elements, but the fact that their relative cosmic abundances fall within narrow windows that support both stable, long-lived stars and complex organic chemistry. This has been called a biophilic tuning of elemental synthesis: conditions that are neither generic nor expected in a random universe.

The Structure Formation Timing Problem

The timeline of cosmic structure formation is finely poised. Observations show that galaxies, stars, and large-scale filaments began forming just early enough in the history of the universe to allow biological evolution to proceed, but not so early as to disrupt the smoothness and expansion of space. If structure formation had begun significantly earlier, gravity could have overpowered the expansion rate, leading to premature collapse or black hole dominance. If it had occurred significantly later, matter would have dispersed too much for galaxies to condense. This requires a delicate balancing act between expansion dynamics, initial density perturbations, and Dark Matter behaviour, none of which are naturally fixed by first principles. The "just right" onset of structure formation is therefore considered a further fine-tuning problem, although the whole context of this discussion is currently being rewritten by the remarkable discoveries of the James Webb Space Telescope (problem #14).

The Matter–Radiation Equality Tuning Problem

In the early universe, radiation dominated the energy density, preventing the growth of structure due to its pressure and relativistic behaviour. As the universe expanded and cooled, there came a moment – matter-radiation equality – when matter began to dominate, allowing density fluctuations to grow into galaxies and clusters. This transition had to occur at just the right time in cosmic history. If matter had come to dominate too early, gravitational clumping would have become too strong, leading to an inhomogeneous, turbulent universe. If it had occurred too late, structures would not have had time to form before Dark Energy accelerated the expansion. The precise timing of this phase transition is not predicted by fundamental physics, but must be tuned by adjusting initial densities of matter and radiation. This introduces yet another layer of unexplained calibration into Λ CDM.

The Amplitude of Primordial Perturbations Problem

The CMB reveals that the early universe contained tiny but nonzero fluctuations in density: about one part in 100,000. These primordial perturbations are critical, for they serve as the seeds from

which all later structure (galaxies, stars, clusters) formed via gravitational amplification.

However, there is a fine-tuning problem in their amplitude. The perturbations had to be:

- Small enough to preserve the smoothness of the CMB and prevent immediate gravitational collapse or black hole formation;
- Large enough to allow gravitational instability to eventually grow them into the vast cosmic web of galaxies and clusters.

Inflationary models can generate such perturbations via quantum fluctuations stretched to macroscopic scales. But the actual amplitude observed ($\sim 10^{-5}$) is not a robust prediction of most inflationary potentials. It must be precisely calibrated by adjusting the energy scale and shape of the inflaton potential – yet another arbitrary tuning.

The Missing Monopoles (#6)

Grand Unified Theories (GUTs) of particle physics predict the production of magnetic monopoles – massive, stable particles carrying a net magnetic charge – during symmetry-breaking transitions in the early universe. According to standard thermodynamic calculations, such monopoles should have been copiously produced in the first fractions of a second. Yet no magnetic monopoles have ever been observed. Their (apparent) complete absence from the observable universe is a mystery: if they exist and were produced as expected, they should now dominate the mass density and (it is presumed) be detectable in cosmic ray experiments. Inflation offers a solution by “diluting” the monopole density through exponential expansion, but we’ve already seen where that leads us.

The Baryon Asymmetry Problem (#7)

A foundational assumption of particle physics and cosmology is that the laws of nature are nearly symmetric between matter and antimatter. In the earliest moments after the Big Bang, the universe should have produced equal quantities of baryons (matter) and antibaryons (antimatter) through high-energy particle interactions. But this is not what we observe. Instead, the universe today is composed almost entirely of matter. Antimatter is exceedingly rare, and no large-scale regions of the cosmos exhibit the annihilation signatures that would result from collisions between matter and antimatter domains. This implies a tiny but decisive imbalance in the early universe – roughly one extra baryon for every billion matter-antimatter pairs. This small excess survived the mutual annihilation that occurred as the universe cooled, and it ultimately seeded all the galaxies, stars, planets, and observers that exist today. The observed baryon-to-photon ratio is approximately $\eta \sim 6 \times 10^{-10}$. This ratio is not predicted by the Standard Model of particle physics. While theoretical mechanisms such as baryogenesis or leptogenesis have been proposed, these require additional conditions beyond known physics, including:

- Baryon number violation (never observed),
- C and CP violation at levels higher than in the Standard Model,
- Departure from thermal equilibrium in the early universe.

Moreover, the existence of observers or life does not straightforwardly explain the asymmetry via anthropic reasoning. A universe with no net baryon asymmetry would likely contain no complex structures at all.

The Hubble Tension (#8)

In recent years, a large and persistent discrepancy has emerged between independent measurements

of the Hubble constant (H_0) – the parameter that describes the rate of cosmic expansion. Resolving this conflict, known as the Hubble tension, is one of the most pressing challenges in contemporary cosmology. It has prompted serious reflection on the assumptions underpinning Λ CDM.

There are two primary and independent methods used to determine the value of the constant, and they yield results that differ well beyond the range of mutual error bar. The first method infers H_0 by analysing temperature fluctuations in the CMB. When interpreted within the Λ CDM model, this method yields a value of 67.4 ± 0.5 km/s/Mpc. This approach is **model-dependent**. It depends on assumptions made within Λ CDM (especially inflation) which do not necessarily apply to other cosmological models.

The second method derives the constant from observations of astronomical objects in the local universe, using the so-called cosmic distance ladder. This process involves calibrating the intrinsic brightness of Cepheid variables. A Cepheid variable is a type of massive star that pulsates in a regular cycle, changing in brightness with a well-defined period. The crucial characteristic of Cepheids is the direct relationship between their pulsation period and their intrinsic brightness (luminosity), a relationship known as the period-luminosity law, discovered by Henrietta Swan Leavitt. This law makes them powerful "standard candles" for measuring vast cosmic distances: by observing a Cepheid's pulsation period, astronomers can determine its true luminosity and then calculate its distance by comparing it to its observed apparent brightness] and Type Ia supernovae⁸. The SH0ES (Supernovae, H_0 for the Equation of State) collaboration, among others, has consistently obtained higher values of 73.0 ± 1.0 km/s/Mpc. This method is relatively model-independent.

The discrepancy between these two values now exceeds 5 standard deviations, which makes it highly unlikely to be attributable to statistical error. While it has been suggested that unrecognised systematic errors may be responsible, extensive reanalyses and cross-checks using different methods and observatories have failed to eliminate the discrepancy. Very recently, the James Webb Space Telescope has essentially eliminated the possibility that the Hubble Tension is just a measurement error in the distance ladder. JWST's high-resolution infrared data has confirmed the Cepheid distances to an unprecedented degree. The tension is now a "Crisis of Physics," not a "Crisis of Data."

The Hubble Tension suggests there is a deep flaw in our understanding of the universe's early conditions, the nature of Dark Energy, or the validity of the Λ CDM model itself. Possibilities under investigation include modifications to the physics of the early universe (such as early dark energy or extra relativistic species), revised models of Dark Matter, and even exotic proposals involving varying fundamental constants or departures from GR. In other words, even more epicycles.

The S8 Tension (#9)

The S8 discrepancy sits at the centre of how we read the universe's structure, and the problem shows up the moment you try to connect the early universe to the late one. The standard model gives you a clear recipe: take the tiny fluctuations recorded in the cosmic microwave background, evolve them forward through billions of years of gravity-driven growth, and you should end up with a very specific level of matter clumpiness today. The parameter S8 is the way cosmologists package that clumpiness, so it becomes a single number that captures both how strong the fluctuations are

⁸ Type Ia supernovae are the thermonuclear explosions of carbon-oxygen white dwarfs in binary systems, occurring when the white dwarf accretes mass from a companion star and reaches a critical limit, the Chandrasekhar limit, triggering an explosion. These explosions are known as "standard candles" due to their consistent intrinsic brightness, allowing astronomers to measure vast cosmic distances and the accelerated expansion of the universe to estimate distances and recession velocities of galaxies.

and how much matter there is to amplify them. Planck pins that number down with remarkable precision, which means the late-time universe is not supposed to have much freedom here.

The trouble is that the universe seems to use that freedom anyway. When you look at weak lensing surveys, at galaxy velocities, or at how many clusters have formed, you keep finding a late-time cosmos that is smoother than the Λ CDM forecast. The shapes of distant galaxies carry the imprint of the mass in front of them, and those distortions consistently point to a lower S_8 . Galaxy clustering offers the same story in a different language, and even the abundance of massive clusters agrees. These are very different observables with very different systematics, yet they all point in the same direction. The early universe says one thing and the late universe says another.

This mismatch isn't a minor calibration wrinkle, because the tension grows as the data improves. Each new survey tightens the low-redshift numbers without drifting toward the CMB prediction, so the two regimes remain stubbornly out of alignment. If the standard picture held together cleanly, the growth of structure would flow smoothly from recombination to today, but instead it bends away from the expected trajectory. Something in that forward evolution is not matching the real cosmos, and the gap has become one of the most persistent and revealing problems in contemporary cosmology.

Dark Energy (#10) and the Cosmological Constant Problem (#11)

Empirical observations, at least when viewed in the context of Λ CDM, suggest that a repulsive force is *pushing the universe apart* at an accelerating rate. Dark energy accounts for roughly 70% of the total energy density in the standard Λ CDM model, but nothing else is known about it.

For most of the 20th century, it was assumed that the expansion of the universe must be *slowing down* due to the mutual gravitational attraction of all matter. Cosmologists expected that measurements of distant galaxies would reveal a deceleration, indicating that the rate of expansion had been higher in the past and was gradually tapering off. However, in 1998, two independent teams published results based on observations of Type Ia supernovae at high redshift that defied this expectation. These standard candles were fainter than expected, suggesting that the expansion of the universe is *accelerating*. This astonishing result required a major revision of the cosmological model: cosmologists revived Einstein's cosmological constant. In modern terms, it corresponds to a constant energy density filling space homogeneously, with a negative pressure that drives accelerated expansion according to the Friedmann equations. Mathematically, the cosmological constant exerts a pressure which causes the expansion of the universe to speed up rather than slow down. Reintroducing this term allowed cosmologists to fit the supernova data within the framework of GR, without altering its fundamental structure.

At least by the lights of the standard model, further confirmation of accelerated expansion came from independent sources. Measurements of the CMB revealed a universe that is spatially flat, yet the total matter content accounts for only about 30% of the critical density (the amount needed for flatness under general relativity). Λ CDM resolves this by introducing Dark Energy to make up the remaining 70%, ensuring consistency with the observed structure of the CMB acoustic peaks, the large-scale galaxy distribution, and baryon acoustic oscillations.

While Dark Energy plugs the particular hole it was invented to plug, it immediately raises profound conceptual problems. Most serious of all is a staggering mismatch between theoretical prediction and observational measurement that calls into question our understanding of both quantum field theory and gravitation. It is not merely a technical inconsistency, but a foundational conflict that has resisted solution for decades. In Einstein's field equations of GR, the cosmological constant Λ appears as a term that counteracts the attractive force of gravity and drives the accelerated expansion of spacetime. Observations over the past few decades have converged on the conclusion that the universe's expansion is indeed accelerating, and that this acceleration can be accurately modelled by including a small, positive value of Λ . In natural units, the observed value of the cosmological constant is approximately: $\Lambda_{\text{obs}} \sim (10^{-3} \text{ eV})^4$. This value corresponds to an

energy density of about $\rho_{\Lambda} \sim 10^{-47} \text{ GeV}^4$, which is small (but nonzero).

From the perspective of quantum field theory (QFT), the vacuum is not empty. Instead, it teems with virtual particles and fluctuating fields, whose zero-point energies should contribute to the vacuum energy density. If one naively sums the zero-point energies of all known quantum fields up to the Planck scale – where quantum gravity effects are expected to become significant – one obtains a predicted vacuum energy of the order: $\rho_{\text{vacQFT}} \sim M_{\text{Pl}}^4 \sim (10^{19} \text{ GeV})^4 = 10^{76} \text{ GeV}^4$. Even with more conservative cutoffs at the electroweak or QCD scale, the predicted value remains much too large. The mismatch between the theoretical prediction and the observed value is often quoted as being between *60 and 120 orders of magnitude*, depending on the energy scale at which the calculation is cut off. **This is by far the largest known discrepancy between theory and observation in the history of physics.** It is not like being off by a few percent; it is like predicting something to be the size of a bacteria and then finding out it is larger than the observable universe.

The cosmological constant problem is not simply that the predicted and observed values are different. It is that there is no known symmetry or mechanism in QFT or GR that would cancel or suppress the vacuum energy to the observed tiny value, apart from extreme fine-tuning. In order for the total cosmological constant to match observations, one must postulate a bare gravitational constant that nearly cancels the enormous vacuum energy. This cancellation must occur to between 60 and 120 decimal places, with no physical explanation for such a precise tuning.

Efforts to solve the cosmological constant problem have ranged from supersymmetry (which cancels bosonic and fermionic contributions), to dynamical Dark Energy models (e.g. quintessence), to modifications of gravity at large scales, to anthropic reasoning within the string theory landscape and multiverse frameworks. Despite decades of work, no consensus has emerged. The cosmological constant problem lies at the intersection of quantum theory, gravitation, and cosmology. Its persistence signals a fundamental incompleteness in our current understanding of the vacuum, the nature of spacetime, and the interface between microphysics and cosmological structure. The problem and its resolution are widely regarded as a key to progress in unifying quantum mechanics and relativity.

Dark Matter (#12)

Dark Matter has never been directly detected, but regardless of that it is now thought to comprise approximately 85% of the matter content of the universe and about 27% of its total energy density. The hypothesis of Dark Matter was not introduced for a single reason, but rather emerged as a unifying explanation for multiple independent observational anomalies across different astrophysical and cosmological scales. In each case, visible (baryonic) matter alone proved insufficient to account for the observed gravitational effects.

1. Galaxy Rotation Curves

The original and most famous evidence for Dark Matter came from the study of spiral galaxy rotation curves. According to Newtonian dynamics, the rotational velocity $v(r)$ of stars orbiting at a distance r from the galactic centre should decrease with distance once outside the bulk of the visible mass, roughly following: $v(r) \propto 1/\sqrt{r}$. However, beginning with the work of Vera Rubin and others in the 1970s, it was found that rotation curves tend to flatten at large radii: stars and gas far from the galactic centre orbit at roughly constant velocities, rather than slowing down. This observation suggests the presence of an extended, invisible halo of mass surrounding each galaxy, whose gravitational influence maintains the high orbital speeds. The discrepancy between the mass inferred from starlight and the mass required to explain the rotation curves is substantial – typically an order of magnitude or more.

2. Galaxy Cluster Dynamics

Earlier still, in the 1930s, Fritz Zwicky observed that galaxies in the Coma Cluster were moving too rapidly to be gravitationally bound if the cluster contained only the mass visible in stars. Applying the virial theorem to estimate the total mass required to keep the cluster from dispersing, he found that the luminous matter fell short by a factor of up to 100. This mass discrepancy in galaxy clusters was later confirmed through X-ray observations of hot intracluster gas (which itself requires deep gravitational wells to remain bound) and gravitational lensing studies showing that much more mass is present than can be accounted for by visible matter.

3. Gravitational Lensing

GR predicts that massive objects curve spacetime and thus bend the paths of light – a phenomenon known as gravitational lensing. When distant galaxies or quasars are viewed through massive intervening structures like galaxy clusters, the degree of lensing observed allows cosmologists to infer the total mass along the line of sight. In many such cases, especially with strong and weak lensing maps, the lensing mass significantly exceeds the luminous mass, reinforcing the existence of large quantities of invisible mass. Importantly, gravitational lensing provides a *direct* measure of total mass, independent of dynamical assumptions.

4. The Bullet Cluster and Analogous Collisions

One of the most striking pieces of evidence comes from observations of colliding galaxy clusters, such as the Bullet Cluster (1E 0657-56). In these systems, the visible baryonic matter slows and interacts during the collision, while the gravitational mass, inferred from lensing, appears to pass through relatively undisturbed. The spatial offset between the baryonic mass and the total gravitational mass strongly suggests the presence of non-collisional mass, consistent with Dark Matter that interacts gravitationally but not electromagnetically. Similar signatures have been found in other merging clusters. This is the strongest evidence against Modified Newtonian Dynamics (MoND). Dark Matter is a necessary placeholder for a real gravitational effect that MoND cannot explain.

5. Large-Scale Structure Formation

Another key motivation for Dark Matter arises from the need to explain the formation of cosmic structure: the growth of density fluctuations into galaxies, clusters, and filaments in the early universe. The standard model of cosmology assumes that the tiny fluctuations observed in the CMB grew over billions of years into the structures we observe today. However, calculations show that baryonic matter alone, coupled to radiation before recombination, cannot grow fast enough to account for the observed structure, especially on small scales. Dark Matter (being non-baryonic and non-interacting with radiation) can begin clumping earlier, seeding gravitational wells into which baryons later fall. Simulations of structure formation match observations only when Dark Matter is included.

6. Cosmic Microwave Background Anisotropies

Precision measurements of the CMB have revealed tiny fluctuations in temperature across the sky, corresponding to density variations in the early universe. The detailed angular power spectrum of these anisotropies depends sensitively on the composition of the universe. The best-fit models to CMB data require a significant component of cold, non-baryonic Dark Matter to reproduce the relative heights and positions of the acoustic peaks. This result is independent of galaxy dynamics and provides a cosmological-scale confirmation of Dark Matter.

In summary

Despite its success in explaining these phenomena within the Λ CDM framework, the true nature of Dark Matter remains unknown. Candidates range from weakly interacting massive particles (WIMPs) to axions, sterile neutrinos, and more exotic possibilities. Decades of experiments have yet to yield definitive evidence for its identity.

The Quantum Gravity Problem (#13)

A central goal of theoretical physics for nearly a century has been the unification of quantum mechanics and general relativity, but the two most successful theoretical frameworks remain conceptually incompatible. Quantum Field Theory has successfully described the electromagnetic, weak, and strong nuclear forces within the Standard Model, but attempts to quantise gravity using the same techniques have consistently run into intractable mathematical and conceptual problems, suggesting a deep structural mismatch between the quantum and gravitational domains. The core difficulty stems from the non-renormalisability of gravity when treated as a quantum field. Unlike other forces, the graviton (the hypothetical quantum of the gravitational field) gives rise to infinite quantities in loop calculations that cannot be systematically tamed using standard renormalisation procedures. This failure implies that GR, when naively quantised, loses predictive power at high energies or small distances – precisely where a quantum theory of gravity is most needed, such as near black hole singularities or in the early universe.

Several alternative approaches have been developed in response:

- **String theory** posits that the fundamental entities are not point particles but one-dimensional "strings," whose vibrational modes include the graviton. This approach achieves a finite framework that includes gravity, but it requires additional dimensions, supersymmetry, and a vast landscape of possible vacua, many of which are physically untestable.
- **Loop quantum gravity** takes a background-independent approach, attempting to quantise spacetime geometry itself. While conceptually appealing in preserving diffeomorphism invariance, it has yet to deliver a full low-energy limit that recovers classical gravity and quantum field theory in flat spacetime.
- **Asymptotic safety, causal dynamical triangulations, and emergent gravity** models also offer alternatives, but none have yet yielded a universally accepted or empirically confirmed quantum theory of gravity.

This persistent failure to quantise gravity raises the possibility that the gravitational field may not be fundamentally quantum in nature. Some have suggested that gravity may be emergent from deeper, possibly informational or thermodynamic principles, rather than a force to be quantised in the conventional sense. Others propose that QM itself may need revision in order to accommodate a fully relational or background-independent theory of spacetime. At stake is the coherence of our metaphysical picture of reality. If gravity fundamentally resists quantisation, this may indicate that the unification of physics cannot be achieved solely through the tools of 20th-century quantum theory. It may instead require a paradigm shift that reconceptualises either the quantum, the gravitational, or both, as emergent from a deeper substratum.

The Early Galaxy Formation Problem (#14)

Within its first year of operation, the James Webb Space Telescope detected massive, well-formed well-formed galaxies at redshifts as high as 13 – meaning they already existed only 325 million years after the Big Bang. According to Λ CDM, such galaxies should not have had time to

accumulate sufficient mass, should not appear so evolved (with defined structure and mature stellar populations), and should be rarer and smaller at such high redshifts. The abundance, size, and apparent maturity of these early galaxies challenges both the timeline and mechanisms assumed in Λ CDM. While some argue that adjustments to star formation efficiency or feedback processes might resolve this, others view it as a more serious anomaly, possibly requiring rethinking cosmic expansion history, matter content, or the nature of early gravitational collapse.

The Fermi Paradox (#15)

The Fermi Paradox arises from a striking contradiction between two widely held premises:

1. **High probability estimates for extraterrestrial civilisations:** Given the vastness of the universe (13.8 billion years old, with roughly 100 billion stars in our Milky Way alone, and likely many Earth-like planets) there has been ample time for life to originate, evolve, and spread across the cosmos.
2. **Complete lack of evidence for extraterrestrial intelligence:** Despite the above, we observe no clear signs of alien life – no signals, no probes, no megastructures, no evidence of interstellar colonisation – not even any promising sign of the most primitive forms of life. Meanwhile, many of us, with the mainstream media taking the lead, clutch furiously at any passing straw: “New planet discovered that is rocky and in the right place for liquid water. Scientists say it might be home to alien life!”

The paradox takes its name from a casual 1950 remark by physicist Enrico Fermi: “Where is everybody?” Given even conservative assumptions about the likelihood and longevity of advanced civilizations, many should have emerged and become detectable by now. Yet, the observable universe remains silent. The Drake Equation – an attempt to estimate the number of communicative civilizations in the galaxy – reinforces the puzzle. Even with pessimistic parameters, it suggests we should not be alone.

Proposed solutions to the Fermi Paradox include:

1. **We’re rare or first:** Life might be extraordinarily rare or unique to Earth (the Rare Earth Hypothesis). Intelligent life could be difficult to evolve, short-lived, or we may be the first advanced civilisation in our galaxy.
2. **They’re not interested:** Advanced civilisations may choose not to colonise or contact less developed species. They might communicate in ways beyond our detection capabilities or could be deliberately observing us while remaining hidden.
3. **The simulation hypothesis:** We could be living in a simulated reality where alien civilisations do not exist or are intentionally excluded.
4. **They’re gone:** Civilisations might tend to self-destruct or succumb to cosmic catastrophes before achieving interstellar expansion.
5. **We’re not looking right:** Our search methods might be limited or misguided—targeting the wrong wavelengths, timescales, or signals. Alien technology might be indistinguishable from natural phenomena.
6. **They’re here, but hidden or unrecognised:** Some suggest controversial evidence such as UFOs or UAPs could be alien probes or visitors that we fail to identify properly.
7. **The Dark Forest hypothesis:** Inspired by Liu Cixin’s novel *The Dark Forest*, this hypothesis portrays the galaxy as a dangerous wilderness where every civilization is a silent hunter. Revealing one’s presence risks annihilation because it is impossible to know another civilisation’s intentions, civilisations can quickly become existential threats, and assuming

goodwill when it is absent can be fatal. Mutual silence, therefore, is a strategy for survival.

As with many other foundational problems explored in this book, the Fermi Paradox is not a problem lacking in proposed solutions; the real problem is that none of the existing solutions command consensus.

The Black Hole Information Paradox (#16)

In the 1970s, Stephen Hawking discovered that black holes are not perfectly black: quantum effects near the event horizon cause them to emit thermal radiation, now known as Hawking radiation. This radiation is effectively random and featureless, carrying no detailed information about the matter that fell in. If a black hole eventually evaporates completely via this radiation, the information about the initial state seems to be irretrievably lost.

This creates a direct conflict with quantum mechanics, which requires that the evolution of a closed system be unitary, meaning that information must be preserved over time. In Hawking's original formulation, black hole evaporation appeared fundamentally non-unitary, implying that quantum theory would fail under these extreme conditions.

Over the decades, this paradox has been reformulated in several ways. While modern developments – such as the AdS/CFT correspondence, quantum extremal surfaces, and the Page curve – suggest that information can, in principle, be preserved, the problem is not entirely resolved. Even with unitary evaporation, it remains unclear exactly how information is encoded in Hawking radiation, where it resides during evaporation, and how interior and exterior degrees of freedom are related. The paradox, therefore, highlights a deep conceptual tension between the principles of quantum mechanics and the classical description of spacetime, raising fundamental questions about the nature of information, locality, and the ultimate fate of matter in extreme gravitational systems.

The Arrow of Time and the Problem of Now (#17)

The Arrow of Time

Nothing is more mysterious than time – the apparent asymmetry between past and future that pervades human experience, biological processes, and thermodynamic systems. This directional flow of time is so deeply woven into our perception of reality that we take it for granted. Yet at the level of fundamental physics, the situation is paradoxical: the laws governing the microphysical world, such as Newtonian mechanics, Maxwell's equations, GR, and the Schrödinger equation, are time-symmetric. They make no intrinsic distinction between forward and backward temporal evolution. Nonetheless, macroscopic phenomena display a striking temporal asymmetry. This is most evident in the second law of thermodynamics, which asserts that entropy tends to increase over time in closed systems. This thermodynamic arrow aligns with our psychological sense of the passage of time, with the causal structure of events, with the directionality of memory, and with the irreversibility of biological and evolutionary processes. Why does a symmetric microphysical substrate give rise to a manifestly asymmetric macroscopic reality?

The standard scientific response invokes initial conditions. As discussed earlier in this chapter, the observable universe began in a state of extraordinarily low entropy. This special boundary condition might be related to the arrow of time, so perhaps the progression of entropy, causality, and memory, follows naturally (although this just forces us back into asking why the universe's initial entropy was so low in the first place). Moreover, even we accept that the thermodynamic arrow is accounted for by low entropy at the Big Bang, it remains unclear why the subjective experience of time should correlate with the thermodynamic gradient. The questions here

straddle physics, neuroscience, and philosophy of mind: Is the psychological arrow of time reducible to entropy increase, or does it point to something deeper?

Several speculative approaches have been proposed:

- Block universe models (e.g., in relativity) deny the flow of time entirely, treating the past, present, and future as equally real. But this view struggles to account for the phenomenology of change and the salience of the present moment.
- Quantum interpretations, such as the two-time formalism or retrocausal models, attempt to incorporate time asymmetry at a fundamental level, though these remain highly controversial and experimentally indistinct from standard interpretations.
- Information-theoretic or consciousness-based accounts suggest that the arrow of time may be tied to the irreversible processing or integration of information in conscious systems – though such views remain philosophically provocative and scientifically neglected.

As with the low-entropy problem, the arrow of time exposes a deep incompleteness in the current paradigm. Our best physical theories describe time as a dimension, but experience treats it as process.

The Problem of Now

One of the most under-addressed problems in modern physics is the absence of any formal representation of the present moment – what we intuitively call “now.” In both special and general relativity time is treated as a dimension akin to space.

The equations governing the universe’s evolution are time-symmetric and do not single out any preferred moment as “the present.” Instead, the universe is modelled as a four-dimensional block – a spacetime continuum in which all events, past, present, and future, coexist equally. In this “block universe” picture there is no objective flow of time, no moment is singled out as ontologically special, and the passage of time and the sense of “now” are considered emergent or illusory. This sharply contradicts direct human experience, in which the present moment appears privileged: we exist now, not in the past or future. Our conscious experience is temporally localised. We make decisions, experience change, and observe the unfolding of events, all in a way that presupposes a dynamically moving present.

Efforts to reconcile this include relational time in quantum gravity approaches, the moving spotlight theory or presentism in metaphysics and attempts to reconstruct “now” from information processing or entropy gradients. None of these are incorporated into Λ CDM or fundamental physics. As a result, there is no place for now in the official language of the universe, even though conscious observers – the very ones constructing physical theories – experience it in every moment. This disconnect between temporal ontology in physics and phenomenological reality gives rise to what is sometimes called the “Problem of Now,” or the hard problem of temporal existence.

The Memory Stabilisation Problem (#18)

The Memory Stabilisation Problem challenges our most basic intuitions about time, continuity, and the reality of the past, and yet most researchers don’t recognise it as a problem. Neuroscientists typically assume the classical brain is stable enough to store memories in physical structures, without asking whether those structures are metaphysically grounded. Physicists working on quantum foundations usually focus on microscopic systems and steer clear of the complexities of consciousness. Even philosophers of mind rarely formulate the issue with the precision it deserves.

What is taken for granted is this: that the past is fixed, and that memory is simply a retrieval

process from physical traces left by prior events. But in the context of Two-Phase Cosmology, this assumption becomes untenable. The past, like the present and future, only becomes determinate in Phase 2. Only then does metaphysical collapse occur, selecting one consistent history from the superposition of timeless possibilities. Until that point, what we think of as “the past” remains ontologically unresolved. It follows that memory is not just about recording events that have already happened, but a metaphysical act of stabilisation: a way in which consciousness participates in the instantiation of history. Memories appear stable to us because they are part of a locally collapsed region of the background superposition: the portion of the cosmos entangled with a conscious agent. But this collapse does not propagate backward to fix the whole history of the universe. It stabilises only what must be coherent for that agent's experience to persist.

This leads to a deep problem: How can subjective continuity, identity, and meaning, all of which depend on memory, arise from a substrate that is, in its fundamental nature, unresolved? If no classical past is ontologically fixed before collapse, how do stable, retrievable memories emerge at all? The standard scientific explanations are insufficient. Decoherence describes how interference between quantum states is suppressed by environmental entanglement, but it does not explain why a single history is realised. Quantum Darwinism offers a model in which certain information proliferates and becomes redundantly encoded in the environment, giving rise to objective-seeming classical states, but it still presumes that “the environment” is already part of a classically instantiated world.

Quantum Cosmology: where is it?

I have made remarkably few mentions of quantum mechanics in this chapter on cosmology, even though this is the other great pillar of 20th-century physics. This omission reflects not just a historical divergence in scientific development, but a deep conceptual rift that haunts contemporary cosmology.

Like relativity, QM emerged in the early 20th century, but it spoke a radically different language. Where relativity describes smooth, continuous spacetime shaped by mass and energy, QM dealt in discontinuities, probabilities, and the strange behaviour of particles at microscopic scales. The wave-particle duality, the uncertainty principle, superposition, and entanglement painted a picture of nature that was probabilistic, non-local and utterly foreign to the classical intuition behind gravity and cosmic geometry. But despite their shared birth, the two theories developed along parallel but disconnected tracks. QM revolutionised atomic, nuclear, and particle physics, while GR governed astrophysics, black holes, and cosmology. Each theory was extraordinarily successful in its own domain, yet stubbornly incompatible with the other.

This divide is evident in Λ CDM. While the model depends critically on QM in its earliest moments, it quickly reverts to a classical framework. Once the inflationary field decays the evolution of the universe is described almost entirely using GR and classical fluid dynamics, supplemented by particle physics inputs (cross-sections, decay rates) that are themselves treated semi-classically. This methodological division allowed cosmologists to make extraordinary progress without resolving the foundational conflict between quantum theory and gravity. In a sense, the success of Λ CDM has been a trap: it permitted the construction of a comprehensive cosmological model without properly integrating the underlying theories.

Nonetheless, QM is present (quietly) in several key areas of standard cosmology:

- **Inflationary Fluctuations:** The initial seeds of cosmic structure – the slight density perturbations that grew into galaxies and clusters – are believed to arise from quantum fluctuations in the inflaton field, stretched to macroscopic scales by rapid exponential expansion. This process is the most direct bridge between QM and large-scale cosmology.
- **CMB Anisotropies:** The tiny temperature variations in the CMB encode information about

those primordial quantum fluctuations. Their statistical properties (nearly scale-invariant, Gaussian, adiabatic) are predictions of the inflationary quantum vacuum.

- **Primordial Nucleosynthesis:** The synthesis of light elements in the first few minutes of the universe relies on nuclear reaction rates derived from quantum field theory.
- **Vacuum Energy (Λ):** The cosmological constant is typically interpreted as vacuum energy – a concept rooted in quantum field theory. Yet calculations of this energy overshoot observed value reveals a profound mismatch between the quantum view of the vacuum and its gravitational effects.

Despite these roles, QM remains something of a hidden scaffolding in Λ CDM: necessary for certain inputs, but excluded from the overall architecture. The model does not incorporate quantum uncertainty in the evolution of spacetime itself, nor does it offer an account of quantum measurement, decoherence, or the transition from potential to actual cosmic histories. These omissions lie at the heart of the model's growing empirical inadequacies.

The absence of QM in the core of cosmological modelling reflects a deeper problem: *we do not have a quantum theory of gravity*. No one knows how to consistently describe spacetime as a quantum entity, so cosmology has proceeded with a pragmatic truce: use QM for matter and fields, and GR for spacetime. This truce has worked... until it hasn't. As the anomalies mount it becomes increasingly clear that a revolution is needed. We need a model that unites QM with spacetime evolution, not just at the Planck scale, but in a way that informs how we understand observation, measurement, and cosmic history at every level. If Λ CDM is a classical skeleton with a few quantum joints, the next paradigm may require a fully quantum nervous system – one in which observation, information, and probability are not just calculational tools, but ontological ingredients. The missing thread of QM may in fact be the central thread, and its proper integration with cosmology may reveal a cosmos far stranger, but also far more unified, than Λ CDM allows.

Chapter 6: The unfinished quantum revolution

It is no exaggeration to say that the Measurement Problem is the most profound and unresolved question in all of physics. It is a foundational rupture between what quantum theory tells us about the world and what we actually observe. The MP arises because the formalism of quantum mechanics describes physical systems in terms of superpositions (mathematical entities representing multiple, coexisting possibilities), but when we perform a measurement, we always observe a definite outcome. How, when, and why does the transition from indefinite possibility to concrete actuality occur? The standard theory offers no internal mechanism to explain this leap. We are left with a vague and problematic concept – the "collapse of the wavefunction" – that has no physical basis within the theory itself. This strikes at the ontological core of reality: what is the world made of, and what role does observation play in bringing it about? Does the act of observation somehow create reality? Is consciousness involved? Or is there an undiscovered objective, observer-independent process that resolves quantum possibilities into actual events? There is very little agreement about the answers to these questions.

In this chapter, I will trace the origins of the MP, examine its classical and philosophical roots, and survey the various interpretations and attempted solutions. We will see that, far from being an isolated curiosity of microphysics, the MP may hold the key to understanding reality itself: the origins of the cosmos, the emergence of time, and the nature of consciousness.

Physics at the end of the 19th century

At the end of the 19th century there was a major disconnect between science and philosophy. Since the publication of Kant's *Critique of Pure Reason* in 1781 they had been heading in different directions. In philosophy the golden age of German Idealism had culminated in Hegel's grand metaphysical system, while science had been through its own golden age of discovery of the mechanistic workings of the material world. These two ways of understanding reality are at odds with each other. The scientific view was that reality was made of material stuff, the behaviour of which could be completely reduced to mathematical laws. From the point of view of the idealistic philosophers, matter had a subordinate existence within the ultimate realm of mind, and the behaviour of reality was, at least possibly or partially, irreducible to mathematical laws. It looked to many people like both of these ways of understanding reality were almost finished. Hegel presented his system as the culmination of philosophy (a view widely interpreted as implying that the major work of metaphysics was complete), while physicists were hopeful that physics was nearing completion – some loose ends needed tidying up here and there. There was no reason to suspect a dramatic paradigm shift was brewing. Physics and philosophy were not just separated, but had almost lost contact with each other. Scientists had no reason to think about reality in Kantian terms of phenomena and noumena, and this generally caused them no problems.

One exception to this general rule had occurred in the formative days of atomic theory earlier in the 19th century. On this occasion the boundary between physics and metaphysics was tested by a disagreement between scientists over the material existence of atoms. In 1808 British chemist John Dalton had discovered that if each chemical element is given a standard weight, then they always combine in fixed ratios: water is one part oxygen to two parts hydrogen. Dalton took

these constant ratios to be evidence of the actual combination of real physical atoms, and proposed that at the smallest level, the material world is made of atoms. Most scientists agreed with him, but a small minority rejected this claim on the grounds that it went beyond the facts – after all, nobody had ever seen an atom, or any direct evidence of the existence of real atoms. When in 1826 Dalton received the Royal Society's medal of honour from chemist Humphry Davy, Davy warned that the word “atom” should be taken to mean no more than “chemical equivalent”. He recognised Dalton's achievement in purely practical terms: a discovery about how to do science, rather than what reality is made of. French chemist Jean Baptiste Dumas agreed that the word “atom” had no legitimate place in chemistry, on the grounds that it goes beyond experience. German chemist Friedrich August Kekulé claimed that the entire debate belonged to metaphysics, and that the question of whether or not atoms actually exist has nothing to do with chemistry. In fact, the first direct human experience of atoms occurred the following year, when Scottish botanist Robert Brown observed pollen grains in water lurching about as they were randomly bombarded by atoms, but at the time it was assumed that “Brownian motion” had some sort of biological explanation.

Another German chemist – Wilhelm Ostwald – proposed an alternative to the atomic hypothesis, based on the laws of thermodynamics. He claimed that atoms were mathematical fictions, and that the base level of reality was pure energy. Bitter disputes followed between the “atomists” and “energeticists”. Ostwald gave a speech in 1895 with the title *On Overcoming Scientific Materialism*: “We must renounce the hope of representing the physical world by referring natural phenomena to a mechanics of atoms....Our task is not to see the world through a dark and distorted mirror, but directly, so far as the nature of our minds permits. The task of science is to discern relations among realities, i.e. demonstrable and measurable quantities...It is not a search for forces we cannot measure, acting between atoms we cannot observe.”

In the late 19th century science, the viewpoint of the majority of scientists was that everything that existed in the world could be reduced to two sorts of entities: matter (or energy) and fields. Both were assumed by scientists to be *real*. It made no apparent difference whether they were considered to be part of phenomenal reality or noumenal reality. Kant's Transcendental Idealism was philosophy, not science, and physicists were not trying to provide foundations for a science of mind. But note Ostwald's chosen vocabulary: “measure”, “observe”, “the nature of our minds”. Since then, the precise meaning of these words and their relevance to the foundational assumptions of physics have become central concepts in a battle over the nature of reality that is far from over.

There were two physical fields: gravitational and electromagnetic (physics was soon to add two more – the strong nuclear field, which binds atomic nuclei, and the weak nuclear field, which breaks them apart during radioactive decay). In classical physics, fields were understood as continuous distributions of force or energy throughout space, and they have very different ranges. Electric and gravitational forces follow inverse-square laws. Magnetic effects, being part of electromagnetism, have more complex distance dependencies. The three classical forces never reach zero: everything in the universe is attracted to everything else by gravity, even if the effect is infinitesimally small at great distances. The two modern forces only act over the extremely short distances that apply to atomic nuclei, which is why they were unknown until the 20th century.

From the viewpoint of classical physics, only two sorts of laws are needed to explain everything – those that govern the motion of matter, and those that explain the behaviour of fields in terms of matter. These laws are all completely deterministic – if you know everything about the current situation, then, at least in theory, you can know everything about the outcome. At the time, deism was a popular belief – the idea that God created the universe like a piece of cosmic clockwork, set it in motion and then left it to look after itself in a completely deterministic manner. This determinism was defined by Newton's laws of motion and his field of gravity, but the field laws of electricity and magnetism were not discovered until the 1860s, by Scottish physicist James Clerk Maxwell. This discovery would lead to the unravelling of classical physics.

Maxwell's laws combined electricity and magnetism – they were two forces, but reducible to

a single field. Quite unexpectedly, these laws cleared up what had until then been a mystery about the nature of light. Fields are associated with matter – if you shake the matter, then you shake the associated field, and this sends waves radiating away from the location of the shaking, just as waves radiate from a pebble thrown into a pond. Maxwell's laws enabled physicists to calculate the speed that electromagnetic waves travel, and this perfectly matched the speed of light, which had already been measured. This led directly to the conclusion that light must be high-frequency electromagnetic waves, and Maxwell also correctly concluded that there must be electromagnetic waves of other frequencies. Sure enough, Heinrich Hertz discovered radio waves in the late 1880s and in 1895 Wilhelm Conrad Röntgen discovered X-rays. Classical physics appeared to be complete.

Or at least, *almost* complete, for there was a fly in the ointment, known as “the Black Body Radiation Problem” or “the Ultraviolet Catastrophe”. Black objects have no intrinsic colour, but they take on a colour when they are heated up. The colour of iron changes as the temperature rises – red hot, white hot, etc... Physicists wanted to know how to calculate the colour, now known to be an electromagnetic wavelength, from the temperature. This, they presumed, must have something to do with the matter in the black body shaking more violently as it heats up, which everybody assumed must follow Newton's laws (though we now know light is emitted by moving electrons, which don't follow Newton's laws). The problem was that if you do the mathematics, the prediction is that black bodies should glow bright blue regardless of the temperature.

Then in 1900 German physicist Max Planck made a breakthrough. Instead of theoretically allowing matter to vibrate at any frequency, he restricted the *energies* of oscillators to discrete values $E = nhf$, where E is the particle's energy, n is any integer, f is the frequency and h is a constant that is now named after Planck. This rule restricts the particles to a finite set of energies defined by the value of hf . This innovation was not intended by Planck to be representative of reality – it was a trick to make the mathematics simpler, and Planck planned to eventually get rid of it by making the constant zero, so that this finite set was so huge that the matter could once again vibrate at any energy it wanted to, or as near as makes no difference. However, when he set the constant to zero, the bright blue glow returned. This problem had a solution he was not expecting – it turned out that if he set h to one specific value ($6.62607015 \times 10^{-34}$), then the predicted colour perfectly matched experimental values. This constant later became known as the “quantum of action”, since it has the dimension of energy multiplied by time, which is known as “action” in classical physics.

Nobody realised that this was the thin end of a wedge that would soon break classical physics apart. The theory produced exactly the right answer, but nobody could make any sense of it, because it directly contradicted the Newtonian view of reality. Classical physics and Planck's theory were both mathematical, but everything in classical physics is *analogue*, like a vinyl record. It was as if classical physicists had been searching for the last record to complete their collection of long players, and now Planck had found the missing recording, but rather than a record, he had found a compact disc. Nobody realised this was heralding a revolution in physics, because nobody understood the implications. Planck's new theory had solved the black body radiation problem, but could only describe reality *digitally*.

Einstein's four papers in 1905

Albert Einstein (1879-1955) was a German Jew who had taught himself higher mathematics as a teenager. At the age of 13 he was introduced to the work of Kant, who soon became his favourite philosopher. In 1905, when working as a patent clerk in Switzerland, Einstein published four physics papers that, combined with his 1915 paper on general relativity, provided the foundation for 20th century physics.

The first paper explained the ability of light to knock electrons out of metal (the photoelectric effect). By the late 19th century, the wave theory of light was universally accepted.

Einstein's paper was the first hint that it also had particle-like characteristics. Photo-electric experiments had shown that increasing the intensity of a light beam does not increase the amount of energy of each electron knocked out – instead it increases the number of electrons knocked out, while the energy of each remains the same. If, however, you increase the light frequency (which changes its colour along the spectrum from red to blue) then you increase the amount of energy per ejected electron. Einstein's explanation was that light behaves like a shower of particles, each with the energy given by Planck's expression $E=hf$ (an equation that contains Planck's constant quantum of action). Another piece of the puzzle had been found to be digital rather than analogue. It was for this discovery that Einstein was awarded a Nobel prize.

The second paper explained Brownian motion, which was revealed to be not biological at all. Atoms, first hypothesised by Leucippus and Democritus in the 5th century BC, were now accepted by almost everyone to be real, especially after Einstein's explanation and Jean Perrin's later experimental confirmation. The energeticists admitted defeat on this question, but there was a twist coming later in the year.

The third paper eclipsed the first two by introducing a completely new concept of space and time. According to Special Relativity, space and time are not absolute but depend on the velocity of the observer, and the speed of light was a new absolute – the same for all observers. This absolute, unchanging speed was also declared to be a limit that no signal can exceed. This theory did not look much like it belonged to classical physics, but there was no quantum element to it either. In spirit, it owed something to Kant, though Einstein did not frame it that way. Special relativity implied a radically new relationship between space and time, which Hermann Minkowski would soon formalise into four-dimensional spacetime, and four-dimensional space time seemed as different to the material world we experience that it should surely count as noumenal rather than phenomenal. However, unlike Kant, Einstein had found a way to place space and time in a realm beyond experience. Einstein's third paper did not cause any immediate metaphysical crisis in science. His fourth paper discussed a consequence of special relativity: the equivalence of energy and matter according to the equation $E=mc^2$. The energeticists, it turned out, weren't completely wrong after all.

Waves and particles

Waves are unlike particles in three important ways. Firstly, a wave spreads out while a particle is confined to a small area. Secondly, a classical particle follows a single trajectory, whereas a wave can propagate simultaneously along many paths. Thirdly, two waves can pass through each other, while two particles coming together suffer a collision.

In 1923 American physicist Arthur Holly Compton discovered another experimental occurrence of Einstein's "particles of light". Compton shone X-rays onto a target containing loosely bound electrons (famously graphite), and was able to detect both the ejected electron and the particle of light that recoiled in the manner of the cue ball in snooker (the "recoil photon").

Then in 1924 French aristocrat Louis de Broglie submitted a PhD thesis to the Sorbonne that proposed a theory of electron waves and predicted a wave-particle dual nature of matter. His thesis professor was not convinced, but sent a copy to Einstein, who backed the idea, and de Broglie passed his PhD. De Broglie was pointing towards a fundamentally new theory of physics. Classical physics had reduced the world to matter and fields, or particles and waves. Planck, Einstein and Compton had all provided reasons why waves must sometimes be thought of as particles, and now de Broglie was saying particles can sometimes be thought of as waves. At this point nobody had filled in all of the mathematical details, and nobody had a clue how reality could be made of stuff that was simultaneously a wave and a particle. Then in 1925 not one but three quantum theories arrived.

Werner Heisenberg represented a quantum system with a set of matrices, so his theory is called matrix mechanics. In matrix mechanics, physical quantities such as position or momentum are represented by matrices. The diagonal elements correspond to possible measurement outcomes,

while the off-diagonal elements encode transitions or couplings between states. The matrices have diagonal entries that represent the *probability* that the system has that specific value, and off-diagonal entries that represent the strength of the quantum connections between possible values for that attribute. So in this system the momentum, position and other attributes of a quantum entity such as an electron is represented by one of these matrices, rather than by a single number.

Erwin Schrödinger represented a quantum system as a wave form and wrote quantum laws of motion, of the sort that waves must obey, which is known as the Schrödinger equation. His theory is called wave mechanics.

Paul Dirac's theory is harder to explain. Dirac showed that quantum states could be represented as vectors in an abstract space, and that different mathematical formulations (matrix mechanics, wave mechanics) were simply different coordinate representations of the same underlying structure.

Heisenberg: reality is a discrete set of "jumps. Schrödinger: reality is a smooth, continuous field (until measurement). Dirac: reality is an abstract geometric state. The mathematics is clear, but the meaning is a mess. Two things are important to note from all this. The first is that quantum theory is purely mathematical, and the mathematics can be represented in several different ways. The second is that *probability is inherent in this mathematics*.

With the mathematics not only complete, but three versions available, each of which can be used in different situations, it was now possible for other scientists and engineers to start using quantum theory for all sorts of scientific and technological work, which would eventually lead to the development of nuclear weapons and power. But there was still a huge unanswered question about the implications of the theory for our understanding of the nature of reality. How can something be a wave and a particle at the same time? What are we to make of the probabilistic nature of the mathematics?

It is often claimed that unless you have a deep understanding of the mathematics, it is impossible to understand the philosophical relevance of quantum theory. I do not have this level of understanding. However, if the claim was actually true then we should expect our greatest mathematicians to arrive at converging conclusions that would eliminate the mystery, which is very obviously not happening. It follows that the differences in interpretation must be the result of the understanding – or rather the misunderstanding – of something other than the mathematics. The truth is that quantum theory does not, on its own, supply any conclusive answers to the intriguing philosophical questions it poses.

Between 1925 and 1935 Bohr and Einstein debated the meaning of quantum theory, and a patchy, incomplete consensus emerged for an interpretation favoured by Bohr, Heisenberg and Max Born. This became known as the Copenhagen Interpretation (CI) (Bohr's institute was in Copenhagen) and it remains the most popular interpretation today, regardless of its many faults. The CI is non-realist or instrumentalist: it denies that quantum theory describes an underlying reality. You can be forgiven for wondering what the hell that is supposed to mean, and many other physicists objected, Einstein among them. Surely, they said, this goes too far: it is premature to conclude that no future technology could reveal a deeper truth, and all we should say is that, for now, we'll be cautious and sceptical. But Bohr was having none of it. "There is no quantum world," he said, "there is only an abstract quantum description." Heisenberg took a similar view: "The hope that new experiments will lead us back to objective events in time and space is about as well founded as the hope of discovering the end of the world in the unexplored regions of the Arctic." In other words, the CI is a forthright attempt to *shut down* questions about the ultimate nature of reality. "Quantum theory is unbelievably weird, but we've nailed the mathematics and you're never going to get any deeper answers so you might as well stop asking all those awkward questions. *Move along please; there is nothing to see here.*"

The originator of the quantum himself harboured deep metaphysical convictions that ran counter to the emerging instrumentalist consensus. Max Planck, the reluctant revolutionary, came to

believe that the quantum revolution pointed not to a lack of reality, but to a different kind of foundation. In his later years, he argued unequivocally for the primacy of mind: “I regard consciousness as fundamental. I regard matter as derivative from consciousness... Everything that we talk about, everything that we regard as existing, postulates consciousness.” For Planck, the quantum was not a reason to stop asking about reality, but a signpost pointing toward a mental, or ideal, ground of being – a path the Copenhagen school would deliberately choose not to follow.

The questions did not go away. For all the Copenhagenists' vehement denials that quantum theory can tell us anything about the nature of reality, their own interpretation raises very specific questions. There are two parts to the CI. The first is that there is no reality in the absence of observations, and the second is that observation somehow creates reality (in practice, CI treats measurement as the point at which quantum possibilities become definite, though Bohr himself avoided saying that consciousness literally creates reality). The problem is that if that's your theory then it is rather important to explain exactly what “observation” means, and the CI doesn't even try. In the words of physicist Murray Gell-Mann in 1976: “Niels Bohr brainwashed a whole generation of physicists into thinking the job was done fifty years ago.”

The Uncertainty Principle

The Uncertainty Principle was formulated by Heisenberg in 1927. It asserts that there is a fundamental limit to the precision with which certain pairs of physical properties, called *conjugate variables*, can be simultaneously known. The most famous pair of these variables are position and momentum. The more precisely we know the position of a particle, the less precisely we can know its momentum, and vice versa. This is not due to technological limitations. It is a fundamental feature of the quantum world.

The Uncertainty Principle can be thought of as a result of the wave-particle duality described above. A particle's position corresponds to a localised state, while its momentum corresponds to a wavelength (from its wave-like behaviour). Localising the particle more precisely in space makes its momentum (or wavelength) less defined. This principle upends classical mechanics, where, at least in theory, you can measure a particle's position and momentum as precisely as you want. It introduces the idea that at a quantum level, reality isn't fully deterministic. The future states of a system can't be predicted with absolute certainty but only in terms of probabilities.

There is also an uncertainty relation between energy and time that implies that during very short time intervals, large fluctuations in energy can occur, which is often invoked in phenomena like quantum tunnelling or the existence of virtual particles in quantum field theory. (Unlike position and momentum, time is not an operator in quantum mechanics, so the energy–time relation is more subtle, but the physical intuition is correct.)

The Measurement Problem (#19)

The Measurement Problem is the most important metaphysical issue raised by quantum theory. The difference between “observation” and “measurement” is subtle. “Observation” implies a human or other conscious observer, whereas “measurement” implies a measuring device, but from the point of view of physics they play the same role: they are the explanation for when, why and how a set of quantum probabilities becomes a single manifested outcome.

Schrödinger's wave function evolves deterministically according to his wave equation, predicting the system's future states, but since it's a wave, it spreads out in multiple directions simultaneously. Despite this, actual measurements/observations always find the system in a definite state. This means that the act of measurement/observation alters the system in a way not explained by the wave function's evolution. To rephrase Steven Weinberg: If the Schrödinger equation can predict the wave function at any time, and if observers themselves are described by this wave

function, why can't we predict exact measurement outcomes, only probabilities? How do we bridge the gap between quantum reality and our conscious experience of a material world in a definite state? This is the Measurement Problem.

Schrödinger came up with a now famous thought experiment to illustrate the implications for our understanding of reality. A cat's fate is linked to a quantum event – the decay of a radioactive atom. Before observation, the atom – and by extension, the cat – is in a superposition of decayed/undecayed and alive/dead states. Yet, when we open the box and observe its contents, we find the cat either alive or dead, and never in a superposition. When, how and why does it stop being in a superposition? Schrödinger did not believe in dead-and-alive cats. He was highlighting a defect in the CI, which does not provide any answer to this question, because it does not define what an observation is.

It is worth noting that Schrödinger was not a physicalist – his view could be interpreted as either idealism or neutral monism – but he never directly connected this metaphysical belief with quantum mechanics. However, we can join some of the dots. He had first been exposed to mystical philosophy through the works of Arthur Schopenhauer, and became a student of the Upanishads. Informally (in letters and essays) he referred to the claim that Atman (the root of personal consciousness) is identical to Brahman (the ground of all Being) as “the second Schrödinger equation”. He had no obvious reason to specify that the box in his thought experiment contained a conscious animal – it would have worked just as well if instead of a dead-and-alive cat, the box contained a spilled-and-unspilled pot of acid, which both ruins and doesn't ruin a hat (“Schrödinger's hat”). This would have removed consciousness from the situation entirely. Then perhaps we could introduce the conscious cat as a variation on the thought experiment. Did Schrödinger believe consciousness has anything to do with the collapse of the wave function? He did not explicitly say so, but if consciousness is central to reality and quantum theory is our best description of reality then how can they not be connected in some way? He made his views clearer in his 1944 essay *What is Life?*, in which he also anticipated the discovery of DNA. The essay ends with a discussion about determinism, free will and consciousness:

"Let us see whether we cannot draw the correct non-contradictory conclusion from the following two premises: (1) My body functions as a pure mechanism according to Laws of Nature; and (2) Yet I know, by incontrovertible direct experience, that I am directing its motions, of which I foresee the effects, that may be fateful and all-important, in which case I feel and take full responsibility for them. The only possible inference from these two facts is, I think, that I – I in the widest meaning of the word, that is to say, every conscious mind that has ever said or felt 'I' – am the person, if any, who controls the 'motion of the atoms' according to the Laws of Nature."

Pauli, Jung, and the Search for a Deeper Unity

While Bohr, Heisenberg, and Schrödinger debated the ontology of waves and particles, one of quantum theory's sharpest minds was pursuing a more radical synthesis. Wolfgang Pauli (1900-1958) was not only a brilliant physicist (formulating the exclusion principle and predicting the neutrino) but also the community's most feared critic, whose intuition for conceptual flaws was legendary. Dissatisfied with the unresolved dualisms of the Copenhagen Interpretation, Pauli sought a framework that could encompass both the physical and the psychological. His collaboration with the psychologist Carl Jung, beginning in 1932, was no mere diversion. Together, they explored the controversial theory of synchronicity: acausal, meaningful coincidences that link inner psychic states with external events. Pauli saw in this a profound parallel to quantum entanglement and non-locality. Just as separated particles remained correlated without a signal, a dream symbol and a simultaneous external event could be meaningfully connected without causal chain. Both

phenomena, he argued, pointed to a deeper, unitary level of reality beyond our classical categories of space, time, and causality.

Pauli called this underlying reality the *unus mundus* (the one world), a neutral monistic domain in which the physical world of matter and the psychic world of meaning are unified. He came to view the "observer" not as a separate consciousness causing collapse, but as an inseparable participant in a reality where the distinction between subject and object is secondary. The act of measurement was not a magical intervention, but the making-explicit of one aspect of this whole, inevitably creating a "cut" (schnitt) in the process.

In a letter to Markus Fierz in 1947, he wrote: "It would be most satisfactory of all if physics and psyche could be seen as complementary aspects of the same reality." In this, Pauli moved beyond von Neumann's chain of observers and prefigured Wheeler's participatory universe, but with a crucial addition: he provided a psychological dimension to the participation. For Pauli, the archetypal patterns of the collective unconscious and the mathematical symmetries of physics were different reflections of the same archetypal order in the *unus mundus*. His work stands as the first major attempt to build a bridge between quantum reality and the reality of the conscious mind, not through simplistic collapse models, but through a vision of a common, transcendent source. He died in 1958, just as Bell was beginning the work that would prove the non-locality Pauli had intuited, leaving his grand synthesis unfinished. I will have more to say about Pauli and Jung in chapter 16.

John von Neumann

John von Neumann was indisputably the greatest mathematician of the 20th century, despite his life being cut short by cancer when he was just 53. He warrants extended discussion, because I have found when I mention his contribution to quantum mechanics, certain people immediately dismiss him with a comment along the lines of "Sounds like Deepak Chopra. I feel pretty safe ignoring that." As we will see below, a similarly dismissive reaction can come from people in positions of much greater authority, such as the author of a textbook on the metaphysics of quantum mechanics published by Oxford University Press – *Quantum Ontology: A Guide to the Metaphysics of Quantum Mechanics* by Professor Peter J Lewis (2016).

Von Neumann was about as far from a scientific heretic as could be imagined. He came from a wealthy Hungarian Jewish family, and was exceptionally gifted. At six he could divide two eight digit numbers in his head, and converse in ancient Greek. At eight he had mastered calculus. An interest in history prompted him to read through the whole of a 46-volume world history series. He attended the best school in Budapest, and his father hired private tutors to provide extra tuition in areas he displayed an aptitude for. These included the famous mathematician Gábor Szegő, who upon their first meeting was brought to tears by von Neumann's talent for producing instant solutions to complex mathematical problems. At 19 he was having mathematical papers published.

In his 34-year career, von Neumann authored over 150 papers, making numerous major contributions to mathematics, physics, computing and economics, and playing a key role in the development of computers, AI, game theory and the hydrogen bomb. At the time of his death he was America's greatest expert and authority on nuclear weapons, and was the originator of the MAD (mutually assured destruction) strategy to control the arms race. In 1999 the *Financial Times* named him Person of the Century, as representative of the 20th century's ideal that the power of the mind could shape the physical world (the double meaning here is unlikely to have been accidental).

In 1932, von Neumann's definitive mathematical analysis of quantum mechanics was published. *The Mathematical Foundations of Quantum Mechanics* is quantum theory's Bible. In it von Neumann mathematically proved that an "ordinary reality" cannot underlie the mathematical quantum facts. He showed that if electrons are ordinary objects, or constructed of ordinary objects, then the behaviour of these objects contradicts what quantum mechanics is telling us about reality.

Here is Lewis (p64) on von Neumann:

“...the worry is that at the end of the day there may be no good way to decide between [the competing solutions to the MP], so the project of basing our ontology on our best physical theory will prove to be impossible, at least at this level of description.

However, there is another line of thought running through the history of quantum mechanics according to which there is no genuine underdetermination here. The measurement postulate has always sat uncomfortably with the rest of quantum mechanics. Von Neumann first explicitly formulated the measurement postulate as part of quantum mechanics, but he also claims that when the measurement postulate applies is arbitrary to a large extent, suggesting that it is not an objective physical process but something more closely connected to our situation in the world as observers. In that case, maybe the measurement postulate should never have been regarded as part of the physical theory itself; perhaps we should read the textbooks as saying that a system undergoing a measurement *looks as if* its state collapses to an eigenstate.”

So far, so good. But having pointed out von Neumann's claim that measurement is not an objective physical process but something more closely related to our status as observers, he does not go on to say anything about any relationship between consciousness and wave function collapse. Instead, he implies that this might mean that the collapse only “looks like” it happened, which leads him in the next sentence to the Many Worlds Interpretation. Von Neumann gets no other mention in Lewis' book, and the most important contemporary defender of this idea – Henry Stapp – isn't mentioned at all. It appears Lewis is leaving out von Neumann's interpretation completely. The only mention of the theory that consciousness causes collapse anywhere in *a book on quantum ontology* is a claim on page 179 about Eugene Wigner, who defended and extended von Neumann's position in the 1950s. Wigner's view that consciousness causes the collapse, says Lewis, “falls short of being possible” because it “requires a deeply problematic interactionist dualism.” Lewis does not explain in the main text exactly what it is that he thinks is impossible, or why. He relegates it to a footnote: “One problem is that it seems utterly mysterious on this view how the ability to collapse wave functions could have evolved with the evolution of conscious beings, since evolution is a purely physical process. For further discussion see Chalmers *The Conscious Mind* 156 and 356.” This is a striking example of how deeply materialist assumptions shape the discourse: the view is dismissed not because it is incoherent, but because it conflicts with an unexamined, unquestioned belief about evolution being entirely physical. The von Neumann/Wigner/Stapp interpretation has been categorically dismissed as *impossible*, on the grounds it is “utterly mysterious” how to square this idea with the evolution of conscious beings. But does materialistic science actually have any idea how conscious beings evolved? No it does not, as Thomas Nagel had already pointed out four years before the publication of Lewis' book.

Quantum foundations before Bell

John Archibald Wheeler (1911-2008) occupies a pivotal place in the history of quantum foundations because he pushed the Copenhagen intuition to its most radical and unsettling conclusions. His later work represents the beginning of the paradigm shift I am describing in this book. Earlier in his career he was the Ph.D. advisor and mentor of two other key figures in the history of quantum mechanics: Richard Feynman and Hugh Everett.

In the late 1940s Feynman came up with a fourth version of quantum theory, called “sum over histories”. To calculate an electron's fate, Feynman adds up all its possible histories, cancels everything out, and whatever is left represents what will actually happen, *expressed as a pattern of probabilities*. This is a computational reformulation which has no impact on the measurement problem. Feynman's formulation makes the ontology even more mysterious: what does it mean for a particle to ‘take all possible paths’? The sum over paths gives a probability amplitude, but to connect it to definite outcomes, we still need a Born rule interpretation and a collapse or selection

mechanism. The Born rule, named after physicist Max Born, states that if a quantum system is described by a wave function, the probability of finding the system in a particular measurable state is proportional to the square of the magnitude of the probability amplitude for that state. In the context of Feynman's formulation, the process calculates a probability amplitude for a specific history or final state. The Born rule is then the essential step that converts this calculated amplitude into a real-world probability that the outcome will be observed in a measurement. Feynman's interpretation retains unitarity (the mathematical equivalent of a complete, uncollapsed wave function), and doesn't explain how or why a single outcome emerges in a measurement, so the collapse problem remains.

In 1951 David Bohm's book *Quantum Theory* argued “electrons are not things” and the following year Bohm appeared to do what von Neumann had claimed was impossible: he created a model that allowed electrons to be ordinary objects without contradicting quantum theory. Bohm's theory is very strange. It involves something called “pilot waves” – a new sort of physical entity, with its own fundamental field and a new law of motion. Quantum entities “ride” on the pilot wave, which is aware of everything else going on in the universe, including “measurements”, and communicates this to the electron. Now the electron can be a normal electron, and the unordinariness is in the wave. Had the world's greatest mathematician made a mistake?

In 1957 Everett invented a radically new interpretation: Many Worlds. Since measurement devices are no different to anything else in the world, measurement interactions cannot be special. Bohr had to assign a special status to measuring devices, conferring on them a classical-style existence that is not possessed by the entities they are measuring. Von Neumann didn't consider measuring devices special, describing them in terms of possibility waves just like any other sort of matter. But in order to justify this, von Neumann had to make the act of measurement metaphysically special – he had to remove it from the rest of reality. MWI gets rid of the act of measurement altogether, by positing that where von Neumann thinks there is an act of measurement or observation, reality splits into multiple diverging timelines.

By the mid 1960s there were already at least five metaphysical theories competing to be the one true interpretation of quantum theory. The first was the original form of the Copenhagen Interpretation. This was Bohr's view, and maybe still is the default position for physicists if you ask them what quantum mechanics tells us about reality. Electrons have no definite attributes when unmeasured. This position flatly contradicts the Newtonian conception of a material reality made of microscopic entities that are themselves ordinary objects with ordinary properties, but fails to replace it with anything but a mystery. This might work if all you're interested in is the practical application of quantum theory, but it raises major philosophical questions that it makes no attempt to answer. Reality is “out there” – it is local – but at the smallest level it is fuzzily undefinable. But what does “smallest level” mean? This interpretation requires a completely arbitrary “Heisenberg Cut” – a fundamental division between two radically different sorts of reality, without any explanation as to why two different sorts of reality exist, or why the line between them should be in any particular place. The big problem with Bohr's view is that he treats measuring devices differently from everything else in reality: everything is treated as a probability wave except for the measuring device. But why should measuring devices be granted immunity to the quantum laws that apply to everything else? If your answer is that the measuring device should be treated as a quantum system that is measured by another measuring device then you have an infinite regress – what is the *real* measuring device? This is called “von Neumann's paradox of infinite regress”, because von Neumann was obliged to break the infinite regress by postulating the collapse of the wave function.

The second was Heisenberg's view, which had already started drifting towards a more observer-centric stance. Quantum theory, according to the CI, represents the world in two different ways: the observer's experience is expressed in the classical language of actualities, while the unmeasured quantum realm is represented as a superposition of possibilities. Heisenberg suggested

we take these representations literally as a model of the way things really are. Thus, according to Heisenberg, the unmeasured world actually is what quantum theory represents it to be: a superposition of mere possibilities. Heisenberg called them *potentia*: unrealised tendencies for action, awaiting the magic moment of measurement that will grant one of these tendencies a more concrete style of being that we humans experience as actuality. It is worth noting that while this view resembles that of von Neumann in important respects, Heisenberg thought of “measurement” as something more epistemological than ontological – measurement as a change in knowledge rather than in reality.

Von Neumann's theory was being championed by Wigner, who in 1961 came up with another thought experiment, which is now known as “Wigner's Friend”. Wigner asked: *what happens when the observer himself is observed?* Imagine an observer (Wigner's "friend") inside a sealed laboratory, performing a standard quantum measurement, such as observing whether an atom is in a superposition of decayed and not-decayed states. According to the standard CI, before observation the atom exists in a superposition of both states. The friend opens the detector, makes an observation, and records the outcome. From the friend's perspective, the superposition collapses: the atom is definitely decayed or not decayed. However, from Wigner's external perspective, the friend himself, along with the atom, the detector, and the friend's brain, can be described by a global wave function that remains in superposition. Until Wigner opens the laboratory and asks the friend for the result, quantum theory (in its unitary evolution form) insists that the entire lab-friend system exists in a combined superposed state. So *when does the collapse actually occur?* Is it when the friend sees the outcome? Or is it only when Wigner gains knowledge of what his friend has seen? Wigner affirmed von Neumann's view that consciousness plays a special, essential role in collapsing the wave function. When the friend becomes aware of the measurement outcome, reality snaps into place: collapse occurs not by interaction with a macroscopic device, but by the intervention of a conscious observer. He wrote: “*It is not possible to formulate the laws of quantum mechanics in a fully consistent way without reference to consciousness.*” Physics, on this view, is not a detached description of a mind-independent world, but is inherently observer-relative.

In Bohm's theory, the quantum potential is nonlocal, meaning the behaviour of one particle can depend instantaneously on the configuration of the entire system. Quantum attributes reside in “the entire experimental arrangement”... which has to end up meaning the whole of reality, since the whole universe could be implicated in a single measurement. Entanglement plays a central role: the pilot wave evolves in configuration space, so subsystems are generally non-locally correlated. And completing this picture was MWI, which denied collapse and thereby attempted to make the whole concept of “measurement” redundant.

Bell's Theorem and its consequences

In 1964 John Bell set out to answer the question raised by Bohm's theory: had von Neumann made a mistake? He had not. Von Neumann's proof involved the caveat that no theory involving *ordinary objects* combining in reasonable ways is consistent with quantum theory. There is nothing ordinary about pilot waves, The wave is nonlocal: its evolution depends on the entire configuration of the system, so it encodes global information instantaneously, which involves faster-than-light transmission of information, which is prohibited by special relativity.

In his attempt to understand what had gone wrong, Bell made the most important advance since 1925. He came up with a theorem (not a mere theory, but a logical proof that will never be overturned) that demonstrates that any model of reality (not just quantum mechanical, but *any* model at all) must be *non-local*. Bell proved that Bohm's superluminal connections are unavoidable. If you believe objective reality is non-local (as Kant did) then this is no problem at all, because the superluminal connections can exist at the noumenal level – they can be part of non-spatio-temporal “reality as it is in itself” rather than part of reality as it appears to us. In contrast, anybody who believes in a local objective reality and accepts special relativity has got some major rethinking to

do. After Bell, any theory of reality has to either be explicitly non-local, or somehow make the local/non-local distinction irrelevant. This was already known in 1966, but is worth mentioning that the 2022 Nobel Prize in Physics was awarded to Alain Aspect, John Clauser, and Anton Zeilinger for their works on “quantum nonlocality”.

Bell's theorem has implications for the options listed in the previous section. MWI avoids Bell's conclusion by denying that measurements have single outcomes – a move that sidesteps the theorem rather than resolving the underlying conceptual tension. Bohr rejected the idea of an underlying reality, but he still treated measurement outcomes as local events, which is incompatible with Bell's result. It cannot work without breaking Einstein's speed limit, so (assuming Einstein was right), Bell's result renders Bohr's local interpretation untenable. Heisenberg's view survived, provided the noumenal “potentia” are considered to be connected superluminally rather than locally. Bohm's view remained alive for the same reason. A recent Nature paper (2025) challenges key assumptions of Bohmian mechanics, though debate is ongoing⁹ Bell's theorem marks the point at which the old metaphysical scaffolding collapses. After Bell, the question is no longer whether reality is no-local, but what non-locality *means*.

Wheeler's Participatory Universe and “It from Bit”

In the 1970s, Wheeler proposed a “Delayed Choice Experiment” in which the experimenter can seemingly decide after a particle has entered an apparatus whether it will behave like a wave or a particle. This delayed choice suggests that the particle's past behaviour – whether it “went through both slits” like a wave or “chose one path” like a particle – is not fixed until the observer makes a choice in the present. Thus, Wheeler entertained the provocative idea that the present can, in a very real sense, influence the past. He did not claim that the present literally changes a fixed past, but that the past is not fully determinate until the measurement context is specified.

Wheeler also introduced the Participatory Anthropic Principle, which speculates that the entire universe might require observers to exist at all. In this view, reality is not a pre-existing static container in which observers appear; rather, the very fabric of space, time, and matter emerges through acts of observation distributed across the history of the cosmos. He likened this to a self-excited circuit where the universe generates observers who then, in turn, help bring the universe into definable existence. Wheeler tried to dissolve the traditional boundary between epistemology and ontology: knowing is being. The observer is not outside the system, looking in, but is entangled with the universe's very genesis. He generally avoided explicit metaphysics in print, but in lectures and correspondence he pushed these ideas much further, hinting that the universe might be fundamentally relational, informational, and participatory.

Wheeler's "It from Bit" suggests physical reality ("it") arises from information ("bit") – that is, from acts of observation, measurement, or binary distinctions made by observers. Wheeler proposed that:

"Every it – every particle, every field of force, even the spacetime continuum itself – derives its function, its meaning, its very existence entirely... from the answers to yes-or-no questions, binary choices, bits."

In other words, *reality is fundamentally made not of matter or energy, but of information*. The bits here are not literal ones and zeroes in a computer, but any elementary quantum event that can be framed as a binary question, like: Is the photon detected here or there? Spin up or spin down? Every such distinction adds a "bit" to the universe's informational history. Measurement isn't just revealing

⁹ Sharoglavova, V., Puplauskis, M., Mattschas, C. *et al.* Energy–speed relationship of quantum particles challenges Bohmian mechanics. *Nature* **643**, 67–72 (2025).
<https://doi.org/10.1038/s41586-025-09099-4>.

pre-existing properties, but helping bring the properties, or even the fabric of spacetime itself, into being. Wheeler's model suggests that without the registration of information, there is no definite reality. This concept places information, observation, and meaning at the heart of physics, shifting away from the classical view of a material substrate that exists independently. Note that he never explicitly claimed that consciousness is required; he spoke instead of 'observer-participancy,' a broader category that includes measurement interactions as well as conscious observation.

John Wheeler believed that spacetime geometry itself arises from accumulated acts of information-registration. His vision of a universe in which information, observation, and reality are inseparable, marks the first major step toward the new metaphysical framework this book completes.

The modern proliferation of quantum interpretations.

The late 20th century saw a proliferation of interpretations seeking to grapple with the measurement problem and the role of the observer. Several notable figures and conceptual advances emerged in this period, further blurring the boundaries between physics, information, and consciousness.

In the 1980s and 1990s, Bohm developed his Implicate Order and Pilot-Wave Theory, offering a deterministic alternative where particles follow well-defined trajectories guided by a non-local pilot wave. Pilot-wave theory is a physical model; his later Implicate Order is a metaphysical extension that attempts to explain the holistic structure underlying quantum nonlocality. Bohm emphasised an underlying holistic order that is not directly observable, where information plays a crucial organising role. He rejected the idea that consciousness collapses the wavefunction, but he did not treat mind as epiphenomenal. In his later work, mind and matter were both expressions of the deeper implicate order. His metaphysics gestured toward a deeper connection between mind and the cosmos, particularly through his dialogues with philosopher-guru Jiddu Krishnamurti.

At the same time, Roger Penrose began arguing that standard quantum mechanics was incomplete without incorporating gravity. Penrose proposed that gravity destabilises quantum superpositions, producing objective collapses — and that consciousness arises from these collapse events. This was later formalised with Stuart Hameroff in the Orch-OR model. Although distinct from Wheeler's informational approach, Penrose's proposals kept alive the intuition that consciousness and quantum processes are intimately linked, even if the precise mechanism remained elusive. Meanwhile, foundational work on quantum decoherence by Zeh, Zurek, and others clarified how interactions with the environment could suppress interference between quantum states, giving the *appearance* of collapse without invoking consciousness. However, this "decoherence" does not by itself select a unique outcome; it only explains why certain outcomes appear classical to observers.

Throughout these decades, there was also a growing philosophical interest in observer-centric interpretations of quantum mechanics, influenced by the epistemic turn in philosophy of science. Thinkers like Abner Shimony and Bernard d'Espagnat explored notions of "veiled reality" and the epistemic limits imposed by quantum theory, highlighting that the observer's knowledge — or lack thereof — might be essential to the fabric of reality itself. The stage was now set for Henry Stapp: Wheeler had introduced the centrality of information and observation, Bohm had proposed a holistic ontology, Penrose suggested a physical basis for collapse involving consciousness, and decoherence had clarified how classicality emerges from quantum systems without solving the measurement problem outright. Stapp's work can be seen as synthesising these influences, pushing toward a psychophysical interactionist model where conscious intentions actively participate in shaping physical reality via quantum dynamics.

Henry Stapp's extension to von Neumann's interpretation

John von Neumann proposed that the measurement problem involves a chain of interactions,

starting from a physical system and moving through the measuring device to the observer's brain, and finally culminating in the conscious observer's mind. He introduced a dualistic framework, distinguishing between two processes:

- **Process 1 (Collapse):** This refers to the collapse of the wave function, a non-deterministic event that occurs when a measurement is made, turning a superposition of possible states into a definite outcome.
- **Process 2 (Unitary Evolution):** This is the deterministic evolution of the quantum system according to the Schrödinger equation, where the system remains in a superposition of states until a measurement (or collapse) occurs.

Von Neumann did not explicitly claim that consciousness *causes* collapse, but he showed that collapse cannot be located anywhere in the physical chain, leaving the observer's mind as the only remaining candidate. On this view, consciousness plays an essential role in resolving the quantum uncertainty, but von Neumann didn't propose a mechanism for how this might work.

In 2007, Stapp extended von Neumann's framework, giving consciousness a more active and fundamental role in the quantum process. He introduced the idea that *conscious intention* can influence physical outcomes through repeated acts of observation, employing the Quantum Zeno Effect (QZE). This is the idea that frequent measurement can prevent a quantum system from evolving, suggesting that conscious attention or focus can "freeze" certain quantum states and influence their evolution. In Stapp's interpretation, conscious choices, made by the mind, select among possible quantum states and repeated acts of conscious attention can stabilise certain outcomes over others, actively shaping the physical world. This differs from von Neumann, who saw consciousness more passively – simply collapsing the wave function without necessarily influencing the physical outcome with repeated decisions. Stapp argued that mental intention plays an active and causal role in determining which possibilities become reality. Note that he rejects classical dualism; instead, he proposes a psychophysical framework in which mind and brain interact through the structure of quantum measurement. Stapp sees the brain as a quantum system in which Process 1 corresponds to mental effort or the focus of attention, which leads to the selection of possible outcomes, and Process 2 is the deterministic evolution of brain states according to quantum laws. These can be influenced by the choices made during Process 1. Conscious decisions can therefore have a direct influence on the physical state of the brain, extending von Neumann's abstract idea of consciousness collapsing the wave function into a concrete model of mind-brain interaction. This provides a model for free will in a quantum context, which Stapp further explored in *Quantum Theory and Free Will: How Mental Intentions Translate into Bodily Actions* (2017).

The quantum trilemma

The metaphysical interpretations of QM represent competing philosophical responses to the Measurement Problem, which arises from a fundamental discrepancy between the formalism of quantum theory and our empirical experience of the world. Specifically, the problem emerges from the mismatch between the mathematical structure of QM, which governs the evolution of the wave function via the Schrödinger equation (a linear, deterministic process that yields an ever-expanding superposition of possible outcomes), and the apparent collapse of these possibilities into a single, definite outcome upon measurement, as consistently observed in empirical reality. Any interpretation of QM must therefore explain how – or whether – this transition from superposition to actuality occurs. They can be classified into three broad categories, which together constitute what I term **the quantum trilemma**. Each of these approaches offers a distinct resolution to the Measurement Problem, yet each also encounters deep conceptual or empirical difficulties.

(1) Physical / Objective Collapse Theories (OC)

Physical or objective collapse theories propose that the wave function's evolution is not purely unitary, but punctuated by real, physical collapse events. Collapse occurs via a real physical process, usually stochastic, modifying Schrödinger evolution.

- **Copenhagen:** Technically agnostic about whether collapse is physical or epistemic, but often treated as implying a physical collapse upon "measurement" (vaguely defined).
- **Ghirardi-Rimini-Weber (GRW):** Collapse is spontaneous, with a low but non-zero rate, especially for macroscopic systems.
- **Continuous Spontaneous Localisation (CSL):** A refinement of GRW, providing a continuous rather than discrete collapse mechanism.
- **Penrose Objective Reduction (OR):** Collapse is induced by gravitational effects when superpositions involve significant mass/energy differences.
- **Flash Ontology (Tumulka et al.):** Collapse is associated with spontaneous space-time events ("flashes"), a relativistic-friendly modification.

The promise of these models is empirical testability. The problem is that no such collapse mechanism has been observed. GRW/CSL are empirically testable in principle, but the parameter space is large and experiments so far have only ruled out some regions. To date, all experimental evidence remains consistent with unitary (collapse-free) evolution of the wave function. As a result, these theories rely on hypothetical processes that remain both empirically inaccessible and theoretically underdetermined, rendering them speculative add-ons rather than explanatory breakthroughs.

(2) Consciousness-Causes-Collapse Theories (CCC)

Von Neumann's formalism allowed the "cut" between observer and observed to be placed arbitrarily, but concluded that only a conscious mind can complete the measurement process. This explicitly contradicts physicalism, since the observer is only being proposed in the first place because it is presumed to be external to the entire physical system being measured. Without that, it doesn't work as a solution to the problem it was created to solve. However, while CCC bypasses the need for an undefined physical mechanism, a different Pandora's box is opened: *How did measurement occur in a pre-conscious universe?* In answer to this question, proponents typically invoke panpsychism or idealism, but these come with their own unresolved theoretical burdens. Most seriously, they entail that minds can exist in the absence of brains – something which is very hard to square with our incontestable knowledge that damage to brains directly causes corresponding damage to the contents of consciousness. It doesn't switch it off (as general anaesthetics do), but it *degrades* it, and it does so in highly predictable ways. How can this be accounted for if brains aren't necessary for consciousness? The panpsychists and idealists do offer answers to these questions, but they have not convinced the critics.

(3) Many-Worlds Interpretation

MWI denies that the wave function collapses at all. It claims that all possible outcomes of a quantum event are realised, each in a distinct and non-interacting branch of an ever-expanding multiverse. Apparent collapse is an illusion. MWI cleanly avoids the central problems of both PC and CCC, but it implies that our minds are continually splitting, and that there are an infinite number of timelines where strange versions of ourselves act in random and inexplicable ways. The role of probability, which is central to quantum predictions, is also thrown into doubt: without

collapse it becomes difficult to explain why observers should expect Born-rule statistics (objective randomness) at all.

The Logical Exhaustiveness of the Trilemma

These three interpretive strategies *appear* to exhaust the logical space of viable responses to the Measurement Problem. Either the wave function collapses or it does not. If it collapses, the cause must be either internal to the physical system (OC) or external to it (CCC). If it does not, then all outcomes must be realised (MWI). Interpretations that attempt to evade this trilemma leave key explanatory questions unanswered. Examples include:

- **Copenhagen (strictly epistemic readings):** Claims "collapse" is just a change in our knowledge, not in reality, which dodges the question of what the world is like ontologically.
- **QBism (Quantum Bayesianism):** Quantum states are subjective degrees of belief. Collapse is an updating of personal knowledge, not a physical event. QBism denies the problem's existence by reinterpreting what quantum mechanics is *about*.
- **Instrumentalism / Shut Up and Calculate:** Refuses to engage with ontology at all, treating QM as a predictive tool only.
- **Consistent Histories** (Griffiths, Omnès): Provides rules for assigning probabilities to entire histories, but leaves the reality-status of unactualised histories unclear.
- **Decoherence:** Explains apparent classicality, but does not resolve why only one outcome is experienced – merely suppresses interference terms.
- **Bohmian Mechanics / Pilot Wave Theory:** Particles have definite positions guided by a real wave function. Collapse is replaced by particle trajectories, but the unobserved branches remain "real but empty." This blends PC (hidden variable guiding equation) and MWI (since the wavefunction itself doesn't collapse), but fails to resolve the measurement problem because of the ambiguous ontological status of the unobserved branches. What does "real but empty" mean? How is it any different to "real but not real"?
- **Transactional Interpretation (TI):** Wave function is real but consists of offer and confirmation waves propagating forward and backward in time. Collapse is the "handshake" between them, an atemporal process. TI reframes the collapse problem but does not solve it because its handshake mechanism does not explain *how or why* a particular absorber is selected from among multiple candidates. It postulates (rather than derives) that the amplitude of these waves gives the probability of transaction formation, effectively inserting the Born rule by hand. The physical process of selection (the actual transition from a probabilistic field of potential transactions to a single, definite outcome) remains mysterious and unmodelled. And while retrocausality (advanced waves) adds conceptual depth, it still shifts rather than dissolves the core question of the Measurement Problem: Why this outcome, not others?
- **Relational Quantum Mechanics (RQM):** The state of a quantum system is always relative to the observer or measuring system. There is no absolute, observer-independent state of a system. A measurement outcome is simply the establishment of a relation between two systems; there is no universal account of "what happened", only what *happened relative to a given observer*. RQM rejects the need for a universal collapse mechanism, but it does not explain how and why one definite outcome arises in the experience of any given observer or what determines the particular correlation (relation) established in a measurement interaction. The Born rule and definite outcomes are assumed, not derived, and there's also

no clear mechanism for reconciling multiple observers' perspectives into a coherent global reality – even if RQM claims this is unnecessary. RQM reframes the Measurement Problem as a matter of perspective but offers no account of the mechanism by which any specific outcome is selected, even relative to a given observer. It is therefore yet another incomplete interpretive strategy that shifts the metaphysical framing but leaves the central explanatory gap unbridged.

Why the Measurement Problem Remains Fundamentally Unresolved

After a century of debate, we are no closer to finding a solution to the MP which can assemble a consensus of either scientists or philosophers. Quantum mechanics works flawlessly in practice, predicting experimental results with extraordinary precision, but the theory remains profoundly ambiguous about the ontology of measurement itself: how or why a single, definite outcome emerges from a spectrum of possibilities. The standard formalism is silent on when, why, how, or even whether the wave function collapses. Each interpretation attempts to bridge this gap by introducing new assumptions, metaphysical commitments, or even entirely new dynamics not yet confirmed by experiment. The absence of empirical means to discriminate between these interpretations deepens the impasse. Decoherence provides a compelling story about the emergence of classicality, but without an actual mechanism for collapse it leaves the problem half-solved. Objective collapse models propose physical mechanisms for collapse, but these remain speculative and unverified. MWI is *bizarre*. And consciousness-based interpretations shift the puzzle into the murky territory of the mind-matter relationship, typically creating as many problems as they solve.

The ongoing multiplication of options is the final insult. Instead of converging on a shared framework, each new proposal opens further branches of disagreement – on ontology, the role of information, the status of probability, and even the nature of reality itself. The diversity of views reflects a fundamental conceptual impasse: the quantum formalism seems to describe an indeterminate, relational, or probabilistic world, yet our experience is one of definitive, stable outcomes. The Measurement Problem is not a puzzle to be solved within quantum mechanics. It is a sign that the ontology underlying quantum mechanics is incomplete. A new framework must explain not only how outcomes arise, but why reality has the structure it does.

The Preferred Basis Problem (#20)

The Preferred Basis Problem is related to the MP, but harder for a non-mathematician to grasp. In quantum mechanics, any system is described by a wave function, which encodes all its possible states. Mathematically, the wave function can be expressed in terms of different bases – that is, different sets of reference states that define how the system's possibilities are represented. For example, if you have an electron, you can describe its state in terms of position (where the electron might be located), momentum (how fast and in what direction it is moving), or spin (a kind of intrinsic angular momentum). Each of these options corresponds to a different basis. The wave function can be rewritten in any of these bases, and the maths treats all of them as equally valid. Quantum theory is invariant under changes of basis: the underlying Hilbert space structure does not privilege position, momentum, spin, or any other representation. However, in practice, when we observe or measure the system, we don't see all possible bases – we see specific outcomes in one basis or another, depending on what our instruments are designed to measure. For example, a screen detecting an electron always records its position, not a superposition of positions and momenta. This creates a puzzle. Why do measurements reveal definite outcomes in some bases (like position or momentum) but not others? What determines which basis nature "chooses" when a quantum system decoheres or when a measurement outcome becomes definite? If quantum mechanics allows infinite equivalent ways of describing a system, but only certain descriptions become realised and stable in the classical world, then surely there must be some physical mechanism, rule, or

interaction that selects the relevant basis. In decoherence theory, the environmentally selected basis is called the 'pointer basis,' but the mechanism that selects a single outcome within that basis remains unexplained. Without this explanation, we cannot fully understand how a continuous, deterministic wave function evolution (which stays in superposition) gives rise to the definite outcomes we see.

Various interpretations of quantum mechanics attempt to explain or dissolve this problem. Decoherence theory suggests that the system's interaction with the environment naturally selects a basis where interference between possibilities is suppressed (leading to position or momentum being "preferred"), but decoherence alone does not explain why only one outcome is experienced. It just explains the apparent classicality. Some interpretations, like Bohmian Mechanics, explicitly privilege position as the fundamental basis. Others assume that all bases are mathematically equivalent, but decoherence ensures that observers experience branching in the basis that becomes dynamically stable. MWI assumes that all bases exist, but observers' experiences align with the decohered basis structures. In CCC it is sometimes posited that the mind selects the basis in which the wave function collapses.

Resolving this problem is essential for any complete understanding of quantum mechanics, because it directly addresses the bridge between the abstract mathematical description of reality and the concrete, definite world we observe. The Preferred Basis Problem also exposes a deeper mystery: the mathematical structure of quantum theory is far more flexible than the world we observe. Something must explain why one particular mathematical representation becomes the fabric of experience.

The Unreasonable Effectiveness of Mathematics (#21)

In his seminal 1960 essay *"The Unreasonable Effectiveness of Mathematics in the Natural Sciences"*, Eugene Wigner articulated a profound puzzle: why should abstract mathematical structures, developed often without regard for empirical reality, so precisely describe the physical world? Wigner marvelled at this "miracle" of correspondence, suggesting that mathematics has a kind of uncanny predictive power – not merely a descriptive utility – which seems to point beyond mere human invention. For Wigner, this was not merely an epistemic curiosity but an ontological hint.

From a materialist perspective, one might attempt to reduce this fit to evolutionary convenience or cognitive bias: humans developed mathematics because it worked, because it mirrored regularities in our environment. Yet this response only shifts the question one level deeper: why should the natural world itself be structured in a way that admits such precise mathematical description? Why does reality not overflow the limits of formal systems, resisting exact capture by equations, but instead repeatedly yield to them?

Several possible explanations have been offered. The Platonist argues that mathematical forms exist in an eternal realm of their own, and the physical universe mysteriously partakes in these forms. The formalist or conventionalist insists that mathematics is nothing but symbol manipulation, its apparent success an artifact of selective application. The formalist view struggles to explain why equations devised for purely internal mathematical reasons – with no empirical motivation – later turn out to describe physical phenomena with astonishing accuracy. The structural realist proposes that the universe is fundamentally mathematical in nature, and that what we call "matter" is secondary to an underlying relational web of quantifiable structure. Where Platonism posits a separate mathematical realm, (ontic) structural realism collapses the distinction entirely, claiming that physical reality just is mathematical structure.

If mathematics were merely a human invention, its uncanny success in physics would be a miracle. The fact that it works so well suggests that mathematics is revealing something about the structure of reality itself.

Chapter 7: The missing science of consciousness

Consciousness is brain activity.
Consciousness is an emergent property of neural networks.
Consciousness is an emergent property of predictive coding in the brain.
Consciousness is an emergent property of multi-scale brain connectivity.
Consciousness is an emergent property of energy flow and information exchange in neural circuits.
Consciousness is an emergent property of information integration across the universe.
Consciousness is an illusion.
Consciousness is the brain's way of modelling itself.
Consciousness is a computational process.
Consciousness is integrated information.
Consciousness is the electromagnetic field of the brain.
Consciousness is synaptic resonance.
Consciousness is global neuronal workspace activation.
Consciousness is biological feedback loops.
Consciousness is a recursive self-model.
Consciousness is a byproduct of language evolution.
Consciousness is a linguistic/social construction
Consciousness is reflexive awareness.
Consciousness is a simulation.
Consciousness is pure information.
Consciousness is the body's way of integrating sensory data.
Consciousness is the sum of adaptive cognitive modules.
Consciousness is a field (like electromagnetism).
Consciousness is a torsion field effect in spacetime.
Consciousness is matter becoming aware of itself.
Consciousness is a fundamental property of the universe.
Consciousness is an evolutionary adaptation for social coordination.
Consciousness is the Bayesian brain generating probabilistic models of the world.
Consciousness is a neural correlate of attention.
Consciousness is a state-dependent pattern of thalamo-cortical oscillations.
Consciousness is a dynamical system near criticality (edge of chaos).
Consciousness is the functional integration of distributed cortical networks.
Consciousness is an epiphenomenon of metabolic processes in neurons.
Consciousness is a computational bottleneck for serial information processing.
Consciousness is the brain's solution to the credit assignment problem.
Consciousness is the dynamic binding of disparate neural signals into unified perception.
Consciousness is the result of high-dimensional attractor states in neural networks.
Consciousness is the synchronisation of gamma wave activity across brain regions.
Consciousness is a form of virtual reality generated by the brain for action selection.
Consciousness is a constraint-satisfaction process optimising for coherence.
Consciousness is the capacity for meta-cognition and error monitoring.
Consciousness is wavefunction collapse? Of course not! Less of the woo-woo please.

The Hard Problem of Consciousness (#22)

What, exactly, is consciousness?

What does it do?

How, why and when did it evolve?

How is it related to physical reality?

Are jellyfish conscious?

As you can see from the list on the previous page, there is no shortage of suggestions as to what consciousness is – and these are just a selection of the sort which are at least *trying* to sound scientific. This is the biggest unsolved mystery in science, and it is a mystery like no other. Can you think of another example of a thing (if “thing” is the right word), which the majority of people agree is in need of some sort of a scientific explanation, but for which there is no consensus as to whether it really exists, let alone what it is?

The sheer abundant variety of “Consciousness is X” statements reflects a widespread belief that if only we can find the correct physical X then the job will be done, or at the very least we will be able to start the work of putting flesh on what is finally the right set of bones. But while the search for X continues in every nook and cranny people can dream up, far less attention is paid to the precise meaning of the other two words (“consciousness” and “is”).

What sort of statements are these supposed to be? Are they definitions? Are they theories? Or are they something else entirely? This is crucial, for the simple reason that if we do not have a stable and widely accepted *definition* of X then our prospects of arriving at a stable and widely accepted *theory* of X are negligible. If they are intended as definitions then they must be rejected, simply because they do not accurately reflect the way people actually use the word “consciousness”. When we say “consciousness” we don't mean “brain activity”, and that applies to everybody, including physicalists. If we did, then the statement “Consciousness is brain activity” would take on the meaning “Brain activity is brain activity”, which communicates nothing at all, and isn't what is meant by the people who say it. We actually use the word “consciousness” to refer to subjective experiences: our own personal experience of reality, and that of other people and at least some other animals. This requires a *private ostensive definition*: we “mentally point” to our own conscious experiences and associate the word “consciousness” with all of that. Then, if we want to avoid solipsism, we must presume our intuition is correct: that we share a reality with other conscious beings. These statements, therefore, do not work as definitions.

We must therefore presume that they are at least *trying* to be theories, but *superficially resembling* theories does not qualify them as actual theories. That would need much more than just an X – the X itself can't amount to theory unless there is also a detailed explanation what the word “is” is supposed to mean. Since consciousness very obviously isn't identical to brain activity or any of the other X's on this list, they are all just vague hypotheses which don't adequately address the central question. If you are claiming X is Y, but X *prima facie* very different to Y, and your theory consists of nothing more than the single word “is”, then you don't actually have a theory. A theory of consciousness must coherently and unambiguously explain *how consciousness is related to everything else that exists* (and brains in particular). Furthermore, if we want to claim the theory is scientific rather than metaphysical then presumably that must include how it is related to the material world, and I say “presumably” because science and materialism are not the same thing, even though they've been joined at the hip for four centuries.

However, the mere act of asking such a question sets up another major problem: if we're saying that there is a *relationship between consciousness and brain activity*, and that *that relationship isn't one of identity* then we've introduced a concept which is both non-material by definition, and which also refers to something which we have acknowledged the existence of. In other words, if consciousness is something material then what else can it possibly *be* except for brain activity? Materialism is the claim that only material things exist, so if consciousness has a

relationship with brain activity, and the relationship isn't one of identity, then materialism cannot be true. The only relationship between consciousness and brain activity which is consistent with materialism is identity. "Is" *needs* to mean "is identical to", or materialism collapses.

Philosopher David Chalmers has called this "the Hard Problem of Consciousness": the problem of explaining how it is possible for consciousness to "arise" or "emerge" from the physical architecture of a living brain. Chalmers' own formulation of this problem involves what he calls "philosophical zombies" (p-zombies) – imaginary entities which behave exactly like humans at all times, but which experience nothing at all. He argues that the very conceivability of such entities should be telling us that the physical facts are never going to be enough to account for consciousness.

While I agree wholeheartedly with Chalmers' conclusion, I cannot accept his argument. I can't conceive of a p-zombie because by definition they behave like ordinary humans at all times, which means that if you asked one whether it is conscious it would necessarily reply with something like "Of course I am! Why are you even asking me that?" In other words, I can't imagine a zombie that believes it is conscious. It might be convincingly human in some respects ...but so is ChatGPT. However, it would *not* be capable of understanding consciousness or anything that depends upon it, or at least not like a conscious being understands those things. I think it would actually say something like "Consciousness? I have never been able to understand what that word is supposed to mean.", in which case it wouldn't be a p-zombie, because that is not how normal humans talk.

My own formulation of the Hard Problem makes clear that the problem is not just hard but impossible. The only sort of material world we can directly know, and can be absolutely certain exists *in the form it appears to exist*, is the one we are consciously aware of. I am certain that I am directly aware of a material world. I presume it is in some way derived from a real, mind-external objective reality, but I cannot be certain what form *that* world exists in. Could it be non-material? Could it exist in the form of non-local information, for example? We have no grounds for ruling such a thing out. Therefore, from our own perspective, the only *material* world we can be certain about exists within consciousness. The Hard Problem is the impossible task of explaining consciousness in terms of a material world when their relationship in the reality as we experience is the other way around.

This line of reasoning has historically led people to non-materialistic metaphysical positions which grant a foundational role for consciousness in the structure of reality – usually some kind of idealism, dualism or panpsychism, all of which solve the Hard Problem by claiming that consciousness has been present in some form since the beginning. Unfortunately, this just sets up yet another unsolvable problem. Scientists may not be able to agree what consciousness is, but that doesn't prevent us from finding a lot of correlations between the contents of consciousness and specific structures and processes in brains (Chalmers calls these the "easy problems"). In other words, we are justified in believing that brains are *necessary* for consciousness, even if they are insufficient (i.e. even if something else is also necessary). It is not possible to reconcile this with the idea that consciousness has been present since the beginning, unless brains were also there at the beginning, which nobody is proposing. This is one of the primary reasons that none of these non-materialistic metaphysical positions has been able to mount a serious challenge to physicalism. They don't offer a secure foundation for a new science of consciousness, and it is difficult to see how they ever could. All three predate the Scientific Revolution, and if any of them are to sustain a future paradigm shift, then we are in need of a convincing explanation as to what brains are for. Brains are too complex and too biologically expensive to be a mere receiver of a complex signal coming from some unknown source via some unknown mechanism, or if it was acting as a "filter" for even more complex information. If that is the way reality works, then what is cognition for?

Consciousness, therefore, is not merely an unsolved problem for science; it is an *unframed* one – a profound gap in the architecture of science itself. The roots of this omission lie in the

origins of modern science. The methodological revolution of the 17th century succeeded by *excluding consciousness*. Galileo and Descartes both argued that objective reality must be defined in terms of mathematical quantities – length, mass, velocity, and force – leaving subjective experience to philosophy or theology. This strategic division allowed physics to flourish, but it came at a cost: the subjective domain of experience was bracketed off as scientifically intractable. Cartesian dualism created a conceptual firewall between mind and matter, and for centuries the physical sciences developed exclusively on one side of it.

The Even Harder Problem of Consciousness (#23)

I must emphasise that the Hard Problem arises specifically within the context of a materialist or physicalist ontology. For dualists, idealists, and other non-physicalists, consciousness is either fundamental to reality or systematically accounted for (even if this causes other serious problems). In contrast, physicalism must attempt to explain how consciousness arises from an ontology that explicitly excludes it at the outset. In this light, “hard” becomes a euphemism for “impossible” but rejecting physicalism does not automatically offer a viable way forward, because there is no consensus on an alternative. Existing responses to acceptance of the intractability of the Hard Problem can be categorised as follows:

Eliminativism: consciousness doesn't exist

Some philosophers and cognitive scientists have responded to the Hard Problem by denying the existence of consciousness altogether. Eliminative materialists contend that our intuitions about subjective experience are systematically mistaken and that a mature neuroscience will dispense with consciousness as a theoretical entity. While logically consistent, this position is deeply counter-intuitive and self-defeating. It sacrifices the undeniable reality of experience to preserve the coherence of materialism.

Idealism: consciousness is everything

In contrast, idealist theories assert that consciousness is the fundamental substrate of reality, with the physical world emerging from or within it. This position has a long philosophical lineage and has gained renewed attention in recent years, but idealism continues to face significant resistance, largely because it appears to minimise the ontological status of the physical world and invites the problematic implication of disembodied minds: that consciousness can exist independently of any physical substrate, brains included.

Panpsychism: everything is conscious

Panpsychism proposes that consciousness is a ubiquitous feature of the natural world, present to some degree in all matter. Rather than emerging from complex arrangements of non-conscious parts, consciousness is posited as a fundamental property akin to mass or charge. This view attempts to bridge the explanatory gap by denying that consciousness must “emerge” at all. While panpsychism has gained increasing traction as dissatisfaction with physicalism grows, it suffers from its own counter-intuitive implications. Do people really believe that rocks are conscious? The answer is usually that rocks are conscious in a very different way to humans, but most of us cannot bring ourselves to believe that rocks experience anything at all. Panpsychism also struggles to explain how micro-experiences combine to form unified macro-experiences – a version of the Binding Problem, which I will turn to shortly.

Physicalist Emergentism

A final position holds that consciousness “emerges” from sufficiently complex arrangements of matter. This form of emergentism attempts to preserve materialist commitments while acknowledging the novelty of consciousness, but doesn't get us any closer to a coherent explanation. To assert that a radically different ontological category (subjectivity) emerges from physical complexity is to posit a form of naturalistic magic unless the nature of this emergence, and the causal interaction between consciousness and the brain, can be clearly explained. Why does consciousness emerge? Under what conditions? Does it exert causal influence, and if so, how? If it does not, how can the brain know anything about it?

I call this wider problematic the “Even Harder Problem of Consciousness”: concluding that physicalism is incoherent is not enough, because currently nothing else can command a consensus either. And yet it is not that we are entirely lacking in clues. For example, even though we have no idea what it is or what it does, we are perfectly capable of turning it off. All we need to do is figure out how that actually works, and we'll have our answer...

The General Anaesthetic Mechanism Problem (#24)

General anaesthesia presents a profound mystery at the intersection of neuroscience, pharmacology, and the philosophy of mind. Despite over a century of clinical use, the mechanism by which anaesthetics induce loss of consciousness remains poorly understood. This is especially surprising given the chemical diversity of substances that act as general anaesthetics: they include everything from simple inert gases like xenon to larger, structurally complex molecules such as propofol or ketamine. All these compounds, despite their wide variation in structure and interaction profiles, have the capacity to completely suppress consciousness in a controlled and reversible way. Crucially, this suppression occurs without killing neurons or causing permanent damage, without shutting down the brain's electrical activity or basic physiological functions, and typically without disrupting subconscious autonomic processes like breathing (in mild sedation) or heartbeat regulation.

What common feature of brain function is being targeted by such chemically disparate substances to produce such a specific and dramatic phenomenological effect? The fact that consciousness can be suppressed so cleanly and reliably, without any permanent damage, implies that consciousness is functionally delicate (sensitive to precise patterns or states of activity) but also biochemically robust (capable of returning unharmed once the anaesthetic wears off). This challenges assumptions in both neuroscience and cognitive science. Standard accounts typically invoke modulation of GABAergic inhibition, disruption of thalamocortical connectivity, or altered synaptic transmission, but none of these mechanisms fully explains the selectivity, reversibility, and universality of anaesthetic effects across such chemically varied agents.

General anaesthetics therefore offer us a tantalising clue about what consciousness might be, and how it might work, but not, it seems, a big enough clue to lead us to the solutions we're looking for. This reinforces the impression that as things stand, the correct answer is not currently on offer. Something similar applies to all the other problems I will be surveying in this chapter, starting with the singular nature of conscious experiences.

The Binding Problem (#25)

How does the brain integrate information from separate neural processes into a unified, coherent experience? In cognitive neuroscience, it is well-established that different features of a perceptual scene (such as colour, shape, motion, depth, and spatial location) are processed by specialised and anatomically distinct areas of the brain. The visual cortex alone is divided into multiple subregions, each tuned to specific aspects of the visual field. Despite this functional fragmentation, our

conscious experience presents itself as a seamless whole: we see a red apple as red, round, and located there, not as a disjointed collection of independently processed features. What binds disparate neural signals into a single phenomenological object? The problem becomes especially acute when we consider that no central neural hub has been identified that performs this integrative function. Even the so-called “global workspace” theories, while offering a framework for large-scale integration, do not yet explain how binding occurs at the level of individual perceptual objects or moments.

Various theories have attempted to resolve the binding problem:

- Temporal synchrony hypotheses suggest that neurons coding for features of the same object may fire in synchronised rhythms, allowing the brain to group them together. However, empirical support for this mechanism remains inconclusive.
- Attention-based models argue that focused attention acts as a kind of spotlight, binding features within its scope. Yet this raises further questions about how the attentional system itself binds information to select targets in the first place.
- Re-entrant processing theories posit that binding emerges from iterative feedback loops between cortical areas, creating dynamic integration over time. But again, this presupposes a mechanism for coherence that is not yet identified.

The Binding Problem lies at the nexus of subjective unity and objective multiplicity. The brain appears to be a distributed, parallel-processing system with no single control centre, but subjective awareness operates with remarkable cohesion. There is an explanatory gap between third-person functional accounts and first-person phenomenology, and as Chalmers has pointed out, such questions quickly bleed into the Hard Problem.

The Frame Problem (#26)

Originally formulated by McCarthy and Hayes in the 1960s as a problem in formal logic, the Frame Problem has since grown into a broader challenge about how any cognitive system determines what is relevant. It concerns how a cognitive system – whether artificial or biological – determines what matters when something in the world changes. How can an intelligent agent efficiently update its knowledge or make decisions without needing to consider every possible consequence of an action or event? At its core, the Frame Problem is about relevance determination: when a change occurs in the environment, what needs to be re-evaluated, and what can safely be ignored?

The Classic Example

Imagine a robot trying to leave a room. As it approaches the exit, the door unexpectedly swings open. Now the robot must decide:

- Should it update its beliefs about the air pressure in the room?
- Should it recalculate the positions of all air molecules?
- Should it consider the effect on room lighting or temperature?
- Or can it simply note that the door is now open and proceed?

Even this seemingly simple scenario reveals the difficulty: if the robot tries to compute every possible side-effect of every event, it becomes paralysed by combinatorial explosion, but if it ignores too much, it risks missing critical changes that affect its goals. If a system must compute what is relevant, it must also compute what is relevant to that computation, and so on – leading to

an infinite regress unless relevance is somehow given rather than calculated. It needs to strike a delicate balance between comprehensiveness and efficiency, and doing so requires a kind of contextual discernment that has never been reproduced in a machine.

Humans, by contrast, usually handle such situations effortlessly. We intuitively know which details are relevant and which can be safely ignored. When the door opens, we do not pause to reconsider the molecular state of the air. In most situations we intuitively know what to do, to the extent that this ability seems almost trivial, and yet is actually a hallmark not only of human intelligence but also of much less cognitively advanced forms of animal life.

Deeper Implications

The Frame Problem reveals deep cracks in the foundations of symbolic AI, where knowledge and rules are encoded explicitly. It also persists, in subtler form, in machine learning and large language models, which struggle with context shifts, long-range dependencies, and implicit relevance. Even with vast data and computation, these systems frequently fail to distinguish the essential from the incidental, and can be easily misled by small perturbations or ambiguous instructions. This problem suggests that intelligence in animals is not merely about rule-following and pattern recognition, but involves the selective framing of the world – a way of constructing and constraining relevance based on goals, attention, embodiment, and meaning. This touches on unresolved questions in epistemology, perception, and the philosophy of mind. In this way, the Frame Problem serves as a microcosm of the greater challenge of artificial general intelligence. Until machines can frame situations appropriately, they will remain brittle, unreliable, and inferior to human minds. However, in the present context the real mystery is not why machines suffer from the Frame Problem, but why humans and other conscious animals don't.

The Evolution of Consciousness (#27)

Modern biology explains the development of life in terms of variation, selection, and adaptation, but it has yet to offer a coherent account of how or why conscious organisms evolved. It cannot say what consciousness is for, or what it does, in a way that fits cleanly within evolutionary logic. Natural selection acts on function. It explains the emergence of complex traits and structures by showing how they enhanced an organism's chances of survival and reproduction. But consciousness, defined here not as behaviour or information processing but as subjective experience, has no clearly defined function. One can describe the adaptive advantages of perception or decision-making without invoking the felt experience of seeing red or making a choice (so this is directly related to the Hard Problem).

This leads some people to the conclusion that consciousness is an accidental, epiphenomenal by-product of brain activity, entirely lacking in causal power. But if that were the case, then it becomes unclear how evolution could have “selected” for it at all, because evolution does not select for non-functional by-products. Either consciousness has a function and influences behaviour, in which case its causal role must be identified, or it does not, in which case its evolution is inexplicable under the current paradigm. Evolution cannot select for properties that make no difference to behaviour, because natural selection operates entirely through behavioural consequences.

Further complicating the picture is the fact that we cannot observe consciousness in others directly. We infer its existence based on behaviour, but it is not at all clear why such behaviour could be, at least in principle, produced by unconscious systems. There is no empirical test to determine whether the behaviour is accompanied by conscious experience. This leads to a second major problem: we have not been able to define consciousness operationally, in terms that science can measure or model. Even if consciousness is somehow tied to brain complexity, there is no agreed-upon threshold where it is supposed to “turn on” (whatever that means), and no clear

evolutionary lineage. The tools of evolutionary biology do not reveal when consciousness first appeared. There is no principled point in evolutionary history where subjective experience is supposed to appear and no empirical method for identifying such a transition.

This was Thomas Nagel's subject in *Mind and Cosmos*. Nagel's central argument is that consciousness, reason, and value cannot be adequately explained by materialist or mechanistic frameworks alone. He painstakingly demonstrates why subjective consciousness cannot be merely an emergent feature of complex brains, and must somehow be a fundamental aspect of reality that demands its own form of explanation. In his view, the standard evolutionary model treats consciousness as an afterthought – something to be accommodated only once all physicalist assumptions are in place, but which fails to do justice to what consciousness is and how it appears in the world. He concludes that the only credible way to explain the evolution of consciousness naturalistically is in terms of a teleological naturalism in which mind is a basic and irreducible part of nature, not reducible to physical processes and not derivable from them through current scientific methods. In other words, he is saying that the most credible naturalistic explanation for the evolution of consciousness is that somehow it was *destined to happen*. He suggests this teleology must have been governed by currently unknown teleological laws, and that we should embark on a search for more examples of teleological processes in nature, so that we can develop a general theory of natural teleology. Nagel's book provoked a great deal of outrage and criticism, but personally I could not fault his arguments. The book you are reading is a direct response to the challenge he issues at the end of it:

"I would like to extend the boundaries of what is not regarded as unthinkable, in the light of how little we really understand about the world. It would be an advance if the secular theoretical establishment, and the contemporary enlightened culture which it dominates could wean itself of [sic] the materialism and Darwinism of the gaps – to adapt one of its own pejorative tags. I have tried to show that this approach is incapable of providing an adequate account, either constitutive or historical, of our universe.

However, I am certain that my own attempt to explore alternatives is far too unimaginative. An understanding of the universe as basically prone to generate life and mind will probably require a much more radical departure from the familiar forms of naturalistic explanation than I am at present able to conceive."

The cause of the Cambrian Explosion (#28)

What, some readers might ask, is *this* problem doing here? What has the Cambrian Explosion got to do with consciousness? The official scientific answer is: "Nothing, as far as we can tell." And yet if we think about it intuitively rather than scientifically then the answer could scarcely be more obvious. Around 540 million years ago, the fossil record reveals a remarkable and rapid diversification of multicellular animal life. Within a few million years, virtually all major animal phyla appeared, alongside numerous other lineages that turned out to be evolutionary dead ends. This detonation of life was one of the most significant episodes in the history of life on Earth, but its underlying causes remain a subject of intense debate. The suddenness and breadth of morphological innovation observed during the Cambrian Explosion defy straightforward explanation within the standard framework of gradual evolutionary processes. Multiple hypotheses have been proposed to account for this unprecedented diversification:

- Environmental oxygenation: Following the Great Oxygenation Event, a secondary steep rise in atmospheric and oceanic oxygen levels may have reached a threshold necessary to sustain metabolically demanding multicellular organisms.

- Ediacaran seafloor anoxia: Oxygen-deprived conditions in benthic habitats might have forced early life to migrate upwards into more oxygen-rich environments, stimulating ecological diversification.
- Ozone layer formation: The establishment of a protective ozone layer in the upper atmosphere would have allowed organisms to venture onto land and shielded DNA from harmful ultraviolet radiation.
- Termination of “Snowball Earth” glaciations: The end of extensive global glaciations may have opened new ecological niches and stabilised climates, enabling evolutionary experimentation.
- Geochemical changes: Increases in calcium ion concentration in seawater may have facilitated the evolution of mineralized skeletons, contributing to new body plans.
- Mass extinction of Ediacaran fauna: The disappearance of earlier complex life forms may have cleared ecological space for Cambrian animals to diversify.
- Planktonic animal diversity: Expansion in plankton diversity could have altered food webs and nutrient cycling, promoting new evolutionary pathways.
- Symbiotic relationships: The rise of symbiosis may have enabled novel biological complexity and interdependence.
- Hydrothermal vent activity: Shifts in deep-sea vent chemistry and location could have created new habitats and selective pressures.
- Evolutionary arms race: Development of chemical defences or predation strategies may have driven rapid co-evolutionary dynamics.
- Geomagnetic and solar influences: Changes in Earth’s magnetic field or bursts of solar radiation may have increased mutation rates or environmental variability, accelerating evolution.
- Extraterrestrial factors: Hypotheses range from galactic starbursts affecting cosmic radiation to speculative ideas of alien introduction of genetic material.
- Astronomical resonance: Transient alignments within the solar system possibly caused unusual tidal forces and environmental changes.
- Microbial coordination: Emergence of collective intelligence among microbes might have enabled new modes of multicellular organisation.
- Genomic innovations: Intrinsic genetic and developmental shifts may have introduced new “genetic technologies” that facilitated complex body plans.
- Evolution of key traits: The appearance of vision, improved neural capacity, or enhanced mobility could have triggered ecological feedbacks promoting diversification.
- Ecological complexity thresholds: Crossing certain complexity or interaction thresholds might have catalysed rapid evolutionary branching.

The Cambrian Explosion remains an unresolved enigma. No single factor, or combination thereof, has achieved widespread acceptance as the definitive cause. Now, let's see, hmm....*what on Earth* could it have been?

The Problem of Free Will (#29)

The problem of free will is a longstanding dilemma concerned with the apparent conflict between human agency and the causal structure of the universe. At its core lies the question: How can we be genuinely free agents if our actions are the outcome of deterministic and random processes? This issue sits at the intersection of metaphysics, ethics, philosophy of mind, and empirical science, raising doubts about moral responsibility, meaning, dignity, and perhaps even personal identity. The classic formulation of the problem arises from the apparent incompatibility of three claims, each of

which seems independently plausible:

1. Determinism: Every event, including human choices, is the product of prior causes governed by natural laws.
2. Free Will: Human agents sometimes make choices that are not wholly determined by prior states of the world.
3. Moral Responsibility: Individuals are rightly held accountable for their actions.

Together, these claims appear logically inconsistent. If determinism is true, it is unclear how agents can be said to "choose", and indeterministic randomness is equally incapable of grounding responsibility. If my action is caused by randomness, then it is not caused by me, and randomness is no more a foundation for responsibility than determinism. The intuition that we are morally responsible is difficult to reconcile with either possibility.

There are three principal philosophical responses:

1. Compatibilism

According to compatibilists, free will and determinism are not in conflict. Freedom consists in acting according to one's desires and intentions without external coercion, even if those desires are themselves causally determined. This view preserves moral responsibility and practical autonomy within a deterministic framework. However, critics charge that it redefines "freedom" in a way that evades the real question: whether human beings originate their actions in any ultimate sense. Kant called it a "wretched subterfuge". To incompatibilists, compatibilism is just a semantic manoeuvre that bypasses the metaphysical problem rather than resolving it, and thereby causes a vast amount of avoidable confusion.

2. Hard Determinism

Hard determinists accept the reality of determinism and reject the existence of free will. On this view, human beings are biological machines and free will is an illusion – perhaps a psychologically useful one, but not ontologically real. The implications are dire: if individuals cannot help but act as they do, then traditional notions of praise, blame, and justice must be rethought. Some hard determinists advocate replacing retributive systems with therapeutic or corrective approaches. Many of us find such a view deeply unsettling.

3. Libertarianism

Libertarian theories assert that free will is real and therefore determinism must be false. Some actions, on this view, are uncaused or are caused by non-physical agents operating outside the chain of physical causation. Libertarianism preserves moral responsibility by insisting that agents are the ultimate originators of their choices. Its critics accuse it of introducing metaphysical entities that cannot be empirically verified, and of failing to provide a coherent account of how non-physical causation operates.

Empirical science complicates matters further. Classical physics operates on a fully deterministic model whereas QM introduces indeterminacy, but it is hard to see how randomness could confer agency anyway. Neuroscientific studies, such as the well-known experiments by Benjamin Libet, suggest that decisions may be initiated in the brain before subjects become consciously aware of them, raising deep questions about the causal role of conscious will.

Thomas Nagel, as ever, offers a nuanced articulation of the problem. In *The View from Nowhere* (1986), he identifies a deep conflict between the subjective standpoint – the lived experience of choosing freely (the "view from somewhere") – and the objective standpoint (the "view from nowhere"), which views human action as part of a causally determined physical system.

This conflict generates what he calls a problem of autonomy: how can we be morally responsible if our actions are ultimately the result of forces beyond our control? Nagel rejects compatibilism as insufficient, not because it is logically flawed, but because it fails to address the deeper existential intuition that true agency must involve some form of origination. He is equally unsatisfied with libertarianism, which seems to offer no intelligible account of how indeterminacy could ground meaningful control. In the end, he gives up:

“I believe that the compatibility question has not been properly formulated, and that nothing currently on the table resolves it. I cannot even tell whether the truth lies in one of the existing views or in some alternative not yet imagined. My impulse is to say something that I know is not really coherent: that we are somehow responsible for what we do, even though it is ultimately a matter of luck.”

The Problem of Meaning and Value (#30)

Perhaps the most intractable problem on this list is not physical or even biological. It is *existential*. It concerns the status of meaning and value in a scientific universe. In *Mind and Cosmos*, Nagel argued that the prevailing materialist paradigm of reductive physicalism fails to account for the very capacities that make scientific understanding possible in the first place: consciousness, cognition, and the perception of value.

According to standard science, the universe consists of law-governed particles and fields, impersonal and indifferent. But we, who are made of those particles, find ourselves asking:

- Why is there something it is like to be us?
- Why do we experience the world as meaningful?
- Why does reason track truth, and why does truth matter?

For Nagel, these are not “soft” philosophical distractions from the real business of physics – they are central features of reality. The emergence of consciousness, intentionality, and normativity cannot be explained by a worldview that treats only physical facts as ontologically fundamental. If value and meaning are real – if they exist – then they must be part of the natural order, not afterthoughts or illusions. The current scientific picture offers no place for such things. Evolutionary biology attempts to explain moral judgment and aesthetic feeling as adaptive functions; neuroscience reduces rationality to neural computation; cosmology regards the entire history of life and mind as a byproduct of blind initial conditions. In this view, truth, goodness, and beauty are accidents. As Nagel puts it: “It is difficult to make sense of the idea that life is something that could be fully understood by chemistry and physics alone.” This leads to an even deeper dilemma. If reason and value are not part of the fundamental structure of reality, why should we trust them at all? If our beliefs are shaped solely by selection pressures or physical law, what warrants our confidence that they correspond to truth, or that truth itself matters? If our cognitive faculties evolved for survival rather than truth, then their reliability becomes a matter of luck, which undermines the very possibility of scientific knowledge.

The problem of meaning and value thus undermines the coherence of the materialist framework from within. If we are the kinds of beings who can perceive meaning, recognise value, and pursue truth, then any complete account of the universe must include these capacities not as anomalies, but as central explanatory data. Nagel calls for a radical rethinking: a new conception of nature that integrates mind, value, and reason into its core ontology, rather than attempting to explain them away. Meaning, value, and reason are not detachable from consciousness; they are aspects of what consciousness *does*.

Part Three:

Two-Phase Cosmology

Chapter 8: 0|∞

Part Three of this book will provide an integrated solution to our thirty problems. The test that matters is how this model of reality compares to both the best model that materialistic science can offer, and by the idealist and panpsychist alternatives which seek to replace it. Which is the best fit for the *existing* empirical data?

I am going to start with Problem #1, which is a problem not specific to any particular model of reality, but foundational to any kind of cosmological theory. Why does anything exist at all? Why isn't there just nothing? The solution to this problem is forced by logic, but it is a logic that the scientific world has never fully acknowledged. And yet if you accept this logic, it turns the rest of the problem space upside down.

How can something come from nothing? (#1)

How can something come from nothing? The answer is simple, and has been known since antiquity: *Ex nihilo nihil fit* – from nothing, nothing comes. If absolute nothingness had ever been real, there would still be nothing now. The existence of anything at all means that, barring a logic-defying miracle, some kind of eternal ground must underlie reality. That leaves us with two options. One is an eternally complex source such as an Abrahamic God – a pre-existent being who chooses a possible cosmos and wills it into being. The other is an eternally simple source – a condition with no prior structure, no determinate content, but infinite potential. The simplest possible paradox: an Infinite Void.

I have never believed in the intelligent designer God which so many Western theists are convinced is real. By the time I was old enough to have formed a view on such things, I had decided it was about as believable as Father Christmas, and I chose Christmas Day to refuse to attend church ever again. Much has changed about my understanding since then, but the idea of God as a kind of CEO and project engineer of reality has never made sense to me. If such a being actually does exist – a God who *thinks*, designed the cosmos, and makes strategic decisions about the course of human history – then I have a lot of questions to ask about the details of Its decision-making. So for me this is not a tough decision – I start my system with an Infinite Nothingness. I write this as 0|∞. This represents the unity of absolute absence and limitless possibility – the paradoxical ground from which all structure arises.¹⁰

This intuition is not new. Across cultures and millennia, thinkers have returned to the same idea, each time with different names. In Hinduism, starting from around 1500BC, it is the unmanifest Brahman, beyond qualities, from which manifest reality (*prakriti*) unfolds. In Taoism, from 6th century BC, it is *Wuji* – the undifferentiated stillness before Yin and Yang. For Madhyamaka Buddhist philosopher Nāgārjuna (c.150-250AD) it is *Śūnyatā* (emptiness). This is not nothingness in the ordinary sense, but the recognition that all phenomena lack intrinsic essence and arise only through dependent origination. In the West it goes back to Anaximander and the *Apeiron*. Plotinus (204-270) called it *the One* – ineffable and prior to all categories of being or thought. Medieval German mystics called it the *Ungrund* – the groundless abyss that underlies God and creation alike. More recently Nishida Kitarō (1870-1945), of the Kyoto School, wrote about

¹⁰ Please note that I'm not trying to *prove* that God doesn't exist. There's nothing to stop people from believing that the first level of structure built on top of the mathematical foundation is a realm where God or gods exist. However, this model does not need or imply any Gods.

Absolute Nothingness, conceived as a dynamic field that holds together both being and non-being.

These traditions converge on the common insight that the deepest ground of reality is not a determinate object, nor a being among beings, but a paradoxical absence that is also infinite presence. Every chain of explanation must end somewhere. We can end in complexity, positing a pre-existent complex God, or a multiverse machinery already loaded with laws, constants, and mechanisms, but this simply shifts the question. *Where did that complexity come from?* The only other alternative is to end in paradoxical simplicity, by recognising that the final ground cannot itself be explained without contradiction, because any explanation presupposes it. The ground must be both self-sufficient and unconditioned. It cannot be fully stated in positive terms. It is not a gap in our knowledge, nor is it a placeholder for future science, but a recognition that without such a paradox, no coherent explanation is possible. Modern logic and mathematics give us metaphors for this situation. Gödel showed that any sufficiently rich system contains undecidable statements – truths that cannot be proven within the system itself. The Void is the axiom that cannot be derived, yet without it no system can be complete.

The Pythagorean Ensemble

The Void is not an empty space or a physical vacuum. It is a pre-physical background from which all consistent mathematical structures may emerge. Because there are no spatio-temporal constraints yet, $0|\infty$ “contains” all coherent and unchanging mathematical forms – all sets of internally consistent mathematical relationships, which includes the totality of all physically possible universes, histories, and processes *except those that encode conscious organisms* (a point I will explain shortly). This is a strong form of what is usually referred to as mathematical Platonism: any logically coherent structure exists, in a timeless and spaceless way, within the realm of formal possibility. I prefer to call it Pythagoreanism. This is because Plato didn't stop at mathematical structure, but also included a whole range of “forms”, such as the “perfect table” from which all other tables are ultimately derived. There are no tables in the ensemble I'm describing – there are structures that resemble tables, but nothing can actually *be* a table without a conscious being such as a human assigning that specific meaning to it. There's no perfection in the ensemble either, for the same reason.

Within $0|\infty$, every mathematically valid but non-conscious cosmos exists in superposition. Not “in parallel universes” in the physical sense, but as informational structures with complete internal logic. Some correspond to universes with no stars, some to universes with strange physics and some to our own universe, including the entire history of our cosmos from Big Bang to Earth's early biosphere. These are not “happening”. There is no time or change yet, only possibility. They simply *exist* as coherent totalities in the Pythagorean sense.

This is a structural and logical starting point, not a temporal one. In some ways it resembles MWI and the multiverse of all cosmoses. It is structurally identical to Max Tegmark's mathematical multiverse, apart from one thing: MWI and Tegmark's multiverse are both essentially materialistic notions, which have nothing to say about consciousness. Therefore they contain all possible timelines, and all possible cosmoses, including (since no exceptions are made) all possible timelines involving conscious beings. Conscious beings in these models are just “along for the ride”. Our own lives, as we experience them, are nothing more than random paths through a meaningless multiverse. There is no difference between possibility and actuality – all worlds are actual.

So there is a crucial difference in the model I am describing: the Pythagorean ensemble is only a collection of possibilities, not actualities. If all of these possibilities were as real as the reality we actually find ourselves in, then the standard anthropic solution to the fine-tuning problems in cosmology would apply: all possible realities exist, and only in those which contain intelligent beings such as humans does anybody ever ask questions about how they came to exist. This would leave us with no new scope for progress on the other problems on our list, including the mind-splitting problem of MWI.

We are now talking about a very different situation to the one where we are wondering how our particular reality can come from nothing. The entire problem space has been turned on its head, and we are left with a very different question: if only a small proportion of these possible worlds become actual worlds, what determines which of the possibilities are actualised? Instead of needing to explain how reality can be constructed from nothing, according to a set of laws we can't explain the origin of, we need some sort of “mechanism” (which is in quotes because it might not be quite the right word for something which is essentially metaphysical) which selects specific realities from the vast range of all possibilities. Just such a mechanism has been staring us in the face since 1957, but until now nobody has noticed it.

A new solution to the Measurement Problem (#19)

My conviction that zero=infinity goes back to my initial defection from materialism to neutral monism in 2002, but this did not lead me straight to the Two-Phase cosmology. My journey down that path began during a discussion on social media about the idea that consciousness collapses the wave function. It had followed the usual kind of trajectory: plenty of people wrongly insisting that science has conclusively ruled out the involvement of consciousness in wave function collapse (“you need to understand that no true ~~Seotsman~~ physicist takes it seriously any more...”), and a question which regularly comes up in such discussions: if consciousness is needed for wave function collapse, then what collapsed the wave function before consciousness evolved? To those who ask it, this question strikes at the heart of CCC. We have clear empirical justification for believing brains are necessary for consciousness, so if consciousness is necessary for wavefunction collapse then what are the implications for the state of the universe before the first conscious organisms appeared? Are we proposing that quantum processes weren't happening? Because that is just absurd. Are we suggesting that the universe wasn't in any specific state? Was everything just a kind of “quantum soup”? How could *anything* have evolved if that was the case?

This question had drawn the customary response: consciousness was always there. It's fundamental to reality. This answer leads us towards panpsychism, idealism and dualism, and while it does indeed answer the question, it also denies that brains are necessary for consciousness. This lines up with some very old and much revered philosophical systems, which makes it popular with many anti-materialists, but these systems clash with neuroscience. Obviously this is denied by people who believe consciousness is the sole substrate of all reality, but it is because of this clash that none of these ancient metaphysical ideas are capable of sustaining a major paradigm shift in the 21st century. They are steadily attracting more interest as the scientific crisis deepens, but even their most ardent defenders aren't expecting a revolution any time soon.

Then it occurred to me that a very different kind of answer is possible – one which is so elegant and retrospectively obvious that it is astonishing that nobody has proposed it before now. It is the first solution to the Measurement Problem to actually escape from the Quantum Trilemma I described in Chapter 5, rather than merely dodging the questions. If wave function collapse requires consciousness, but currently there isn't any consciousness because evolution has not yet produced the right kind of organisms, then it logically follows that the wave function isn't collapsing at all. This resembles MWI, except for that there is no material reality, because the Everettian branches exist only as possibilities. They are like Heisenberg's “potentia” but explicitly ontic rather than epistemic. They have not been “realised”. In my terminology, the quantum *physical* world exists, but there is no classical *material* reality yet. When I pursued this line of thought, a new kind of cosmological and metaphysical model began to take shape.

Introduction to the Two-Phase solution (2PC)

The first thing to note this system is very specifically a **non-panpsychist neutral monism**: a metaphysical system where mind and matter emerge together from a more fundamental

neutral/informational substrate. It involves two different kinds of physical reality, which I call “phases”, which exist hierarchically within $0|\infty$.

Phase 0: $0|\infty$.

Phase 1: a non-local, non spatio-temporal, multiverse of uncollapsed possibility. Pure mathematical structure. Quantum physical rather than material.

Phase 2: a local reality where, from our perspective, the material universe exists within consciousness and the wavefunction is continually collapsing (the “mechanism” for which I have not yet specified). Classical material rather than physical.

We could say Phase 2 is reducible to phase 1, which in turn is reducible to Phase 0. We might also say Phase 1 to emerges from Phase 0, and Phase 2 emerges from Phase 1. This scheme is *reductive* because there is a nested hierarchy where each level can be analysed in terms of its predecessor. It is also *emergent*, because each phase introduces qualitatively new properties (e.g. locality, consciousness, collapse) that are not present in the prior phase. This duality reflects a dialectical structure: each phase is both a consequence of and a transformation of the one before it. The system is reductive in structure, but emergent in function and experience.

This system emphasises, rather than attempting to gloss over, the two distinct notions of “physical”. There is no arbitrary boundary like the one between “micro world” and “macro world”. The “physical” part of phase 2 is conceptually the same as classical *materialism*: a world of three spatial dimensions which changes as time “flows” (whatever *that* means). It is a Newtonian-Einsteinian concept, where quantum effects are either completely hidden or only hinted at in very specific and unusual situations. Phase 2 is also directly related to *idealism*, since the material world in question appears within consciousness. Phase 1 is neutral/informational, but it is also “physical” in the quantum sense (it is explicitly non-local and superposed). This bifurcates the concept of “physical” into two different things, similar to the Kantian distinction between noumena and phenomena, except in this system noumenal reality is only partially unknowable (because we can never observe what is happening inside Schrödinger's box) rather than completely unknowable and not even cognisable. It therefore also qualifies as a version of scientific realism because it explicitly states that there is such a thing as a mind-external objective reality, and grants that science does tell us *something* about it. We can't tell whether Schrödinger's hat is ruined, unruined, or both at the same time, but we can be absolutely certain that it is indeed a hat, and not an umbrella. We can theoretically calculate the probabilities of all possible contents of the box. What we can't do is have knowledge of which of these states we will observe when we open it.

What we have here is a new interpretation of quantum mechanics, and it offers a radically new way of escaping from the quantum trilemma I described in chapter 6. It is effectively a synthesis of Many Worlds and Consciousness Causes the Collapse – something which at first glance seems unthinkable, because MWI and CCC are thoroughly incompatible. The breakthrough idea is the realisation that while such a combination cannot work simultaneously, a sequential combination opens the door to a theoretical revolution. This retains the best ideas from both, playing to the strengths of each, while simultaneously eliminating their biggest weaknesses. MWI provides the immense “computing power” required to summon the first conscious animal from the Pythagorean ensemble, but mind-splitting is cut off at exactly the moment it would start happening. Meanwhile, CCC's before-consciousness problem is solved without implying disembodied minds or conscious rocks, because there's no collapse until something that qualifies as a brain has evolved. It's not physicalism, so there is no Hard Problem, and a new way of thinking about the Problem of Free Will becomes possible. However, what sealed the deal, for me, was the link between this new solution to the Measurement Problem and Thomas Nagel's radical proposal in *Mind and Cosmos*.

Nagel says almost nothing about quantum mechanics in that book – he goes no further than claiming that quantum indeterminacy provides sufficient scope to allow teleology in the evolutionary process. His main conclusion is that if physicalism is false then the only naturalistic way that consciousness could have evolved is if the universe somehow *conspired to make it happen*. The lightbulb moment was when I realised that this Two-Phase model provides a *structural* explanation for Nagel's proposed teleology, *without requiring any teleological laws*. Because Phase 1 is like MWI, with all possible outcomes existing in parallel branches of a multiverse, it is logically inevitable that a conscious organism will evolve in at least one of them (if it is physically possible, then it happens). At this point MWI will cease to be true, the timeline leading to the evolution of conscious organisms will become real and all the others will remain as eternally unrealised structures within primordial Phase 1. This would produce a realised timeline where absolutely everything required for the evolution of consciousness has actually happened (or so it appears), even if it was exceptionally improbable.

Phase 2 is the reality we are participating in: a classical, material reality that is instantiated *through* consciousness, though not *as* consciousness. This must be stated carefully. I am not proposing a unified “consciousness field” within which a single material world is globally realised. Rather, Phase 1 functions as a non-spatiotemporal informational structure, while Phase 2 consists of locally rendered classical realities, each anchored to a conscious perspective.

A useful analogy is a virtual game world. Phase 1 corresponds to the total informational structure that defines all possible consistent game states, while Phase 2 corresponds to the rendered view available to each player. There is no need for a single, globally rendered game world; coherence is maintained through the shared underlying structure, not through a unified display. Consciousness will play the role of screen and input device, while the material world is the rendered image itself. This perspectival rendering is the deep reason why Phase 2 physics is relativistic in structure while remaining effectively classical in appropriate regimes. Relativity reflects the absence of any privileged global rendering, not the existence of a literal spacetime manifold. Spacetime is not fundamental. It is a mathematical abstraction that would correspond to a globally unified Phase 2 rendering—a structure that does not, in fact, exist. The objectively real domain is Phase 1: non-local, non-spatiotemporal, and quantum. Spacetime survives as an extraordinarily powerful emergent framework, but not as the ontological ground of reality.

I must emphasise that Phase 1 never completely disappears. “Phase 1” and “Phase 2” have two different (but closely related) meanings in this system: two stages in cosmological history and two ontological realms. In historical Phase 1 there is no ontological Phase 2. All that exists is the virtual world (because there aren't any players yet) and its history, from its beginning to the exact moment that something capable of serving as a player's avatar appears in that world. Historical Phase 2 starts when the first player enters the game, which means that an ontological Phase 1 begins to exist in a new, participatory form. It is no longer a static, unchanging structure, because there's now a player embodied in the game, and that player's decisions, including their choice of what to observe, are dynamically altering the phase 1 structure. I call the both phases “historical” if I am talking about cosmological history, and if I am talking about metaphysics then I call the first kind of Phase 1 “primordial” and the second kind “participatory”.

The Unreasonable Effectiveness of Mathematics (#21)

In the framework of Two-Phase Cosmology, the unreasonable effectiveness of mathematics is not simply a human invention nor an alien Platonic realm. Rather, it emerges as the language of possibility within Phase 1, the domain of timeless superposition where all physically consistent structures reside. In that phase, mathematics is not “applied” to reality but constitutes the generative scaffolding of what can exist. When a particular cosmos is drawn into Phase 2 embodiment, the mathematical skeleton of its possibilities is carried forward into actuality.

The reason mathematics so effectively describes reality, then, is not accidental but intrinsic:

reality itself is sculpted out of possibility-space by processes that are mathematical in form. What Wigner experienced as “miracle” is, in this view, the deep continuity between the abstract order of Phase 1 and the embodied realisation of Phase 2. The predictive power of mathematics arises because it does not merely describe nature – it discloses the architectures out of which nature is woven. Because Phase 1 is inherently mathematical, the structures that survive into Phase 2 are those whose mathematical form is stable under actualisation – which is why certain bases, symmetries, and invariants appear privileged in the classical world.

Chapter 9: In search of a threshold

If you have got a two phase model of cosmology and metaphysics then a new critical question arises. What defines the transition from Phase 1 to Phase 2? What qualifies as a brain in this model? What could be so special about a brain that it triggers or enables a fundamental change in the structure of reality?

If this model is correct, then just as life must have emerged from nonlife, conscious life must have first emerged from unconscious life. These two transitions differ in kind. Life is a process, not a thing, and its beginning was a continuum from chemistry to biology, not a single event (although there must have been multiple critical events along the way). The boundary of consciousness, as it has been defined here, cannot be this blurred. You cannot have half a consciousness or be half conscious; if you are “half conscious,” then you are, by definition, conscious. Somewhere in evolutionary history, the first organism crossed that metaphysical line, and for the first time, the universe experienced itself from within.

Why we need a threshold

Before diving into the evolutionary and metaphysical implications of the Embodiment Threshold, it's worth clarifying why the very idea of such a threshold is unavoidable. The existence of a threshold of this kind does not depend on any specific theory of mind or interpretation of quantum mechanics. It follows directly from a simple logical necessity: either consciousness has always existed, or it appeared at some point in the history of the cosmos. There is no coherent middle ground.

If we reject the notion that consciousness is a fundamental property of matter (as in panpsychism) and the reverse (as in idealism), and also reject dualism, we are left with the recognition that at some stage, physical complexity and informational integration must have reached a critical organisation where subjective experience became possible. Crucially, this threshold is not just an arbitrary milestone in evolutionary history; it represents a categorical shift in the nature of reality itself. Below the threshold, physical systems evolve without any intrinsic “view from within.” Above it, the same physical substrate now “carries” (somehow) a dimension of being that cannot be reduced to its prior components. Something new has entered the cosmos: subjectivity.

The concept of such a threshold is already implicit in many scientific and philosophical frameworks, including materialistic/physicalist ones. Neuroscience assumes that consciousness depends on the organisation and dynamics of the nervous system, implying that some configurations support awareness while others do not. Evolutionary biology sometimes assumes that subjective experience must have arisen gradually, yet it also implies there must be some definite point when it emerged. It doesn't spell this out in its own language, but the logic of the discipline forces something like a definite point of emergence once you take subjective experience seriously. Evolution works through continuous variation, but the traits it selects for are still bounded in the sense that an organism either has a capacity or it doesn't. There can be grey zones and partial precursors, but every trait that matters for survival eventually crosses a line where its functional effect becomes real enough for selection to notice. If consciousness is treated as a biological trait, even a strange one, the same logic applies. There can be proto-forms and precursors, but there is still some moment when the system first supports the minimal structure that counts as experience rather than non-experience. The pressure comes from another angle too. Evolutionary explanations

rely on causal stories about why a lineage did better after some change. If subjective experience is part of that story at all, then there must be a point before which it did not play any role and a point after which it started to influence behaviour, learning, or survival. That dividing line doesn't need to be sharp in time, but it has to exist conceptually because without it there is no way to say when consciousness entered the adaptive landscape. Even if someone argues for a gradual emergence, the moment the system crosses from zero to non-zero experience is still a definite transition. It can be tiny and almost invisible from the outside, but it is a real shift because having even the faintest spark of subjectivity is not the same as having none at all. This is how evolutionary thinking ends up implying a threshold whether it wants to or not. Even computational theories of mind concede that there must be a minimum condition – a threshold of complexity or functional integration – before any system could, in principle, host a first-person perspective.

In the system I am describing, the ET is the point at which there is such a thing as *what it is like to be* a physical structure (a living brain in a body, until such time as we figure out how to replicate this function in a machine). There must have been a first conscious animal – not merely a species but a single individual that was the first to arrive at the ET and, in doing so, collapsed the primordial wavefunction. This is the creature I denote LUCAS (the Last Universal Common Ancestor of Subjectivity).

The threshold in contemporary science

In recent decades, several scientific models of consciousness have begun to approach (often unknowingly) the conceptual terrain of the ET. Each seeks to identify the precise physical or informational conditions under which mere computation or neural activity becomes experience. Though their languages differ, they all gesture toward a boundary between systems that only *process* information and systems that *feel* it. Among these, two frameworks are particularly clear and ambitious: Integrated Information Theory (IIT) and Global Neuronal Workspace Theory (GNWT).

Integrated Information Theory begins from the claim that consciousness is identical to the intrinsic causal structure of a system – what the theory calls integrated information, or Φ (Phi). According to IIT, a conscious system is one whose internal parts are so deeply interdependent that the whole cannot be reduced to a collection of separable components. Such a system generates information “for itself,” producing patterns of cause and effect that exist independently of any external observer. When a physical structure reaches a nonzero value of Φ – when it possesses irreducible, self-referential causal power – it has, in principle, crossed from mere mechanism into subjectivity. This represents a quantitative expression of the ET. The moment Φ becomes nonzero, a system acquires an interior dimension – it *exists for itself*. IIT thus offers a mathematical glimpse of what, in metaphysical language, is described as the first local collapse of potentiality into lived perspective.

Global Neuronal Workspace Theory, by contrast, approaches consciousness through functional architecture. It proposes that a system becomes conscious when information, instead of remaining trapped in specialised subsystems, becomes globally available across the whole network. In the human brain, this “workspace” involves wide-scale, recurrent connectivity linking distant cortical regions. Conscious access arises when a particular signal ignites this network, broadcasting its contents to memory, attention, and decision-making systems. In this view, the ET is crossed not when information becomes irreducible, but when it becomes globally integrated and self-sustaining – when the nervous system evolves the long-range loops necessary for information to be *about itself*. The threshold here is architectural: a shift from isolated local processing to a unified field of access and reflection.

Both frameworks, despite their differences, converge on the intuition that consciousness appears when a physical system attains a new level of self-relation. IIT frames this in terms of intrinsic cause-effect closure; GNWT frames it in terms of functional accessibility and global

coordination. However, neither IIT nor GNWT tells us why informational or architectural integration should be accompanied by subjective experience at all. The Hard Problem persists. Informational complexity, regardless of how it is refined, is not equivalent to subjectivity. Neither of them has anything to say about wavefunction collapse either, so there is no integration with the Measurement Problem. In 2PC these models can be reinterpreted as attempts to identify the minimal physical architecture capable of participating in collapse: the point at which a system becomes an ‘observer’ in the technical, not anthropomorphic, sense.

The Quantum Zeno Effect

The account above establishes the ET as the categorical moment when a living system acquires a first-person perspective and, in doing so, enacts a local metaphysical collapse. This section situates that thesis in relation to an influential contemporary voice who came to a very similar structural conclusion from within quantum physics: Henry Stapp. Because Stapp is often misread or caricatured, it is important to say precisely what he achieved and where his theory stops.

Stapp proposed a psychophysical parallelism in which each conscious event maps onto a quantum measurement event. His technical language uses von Neumann’s distinction between unitary evolution (Process 2) and reductions associated with observation (Process 1, and the stochastic outcome that follows). Stapp’s move is to identify the experiential event with the reduction itself: the coming-into-being of an answer from the pool of quantum possibilities is at once a physical transition in the brain and the subjective moment of awareness. In his terms, consciousness and collapse are two facets of the same single process.

This identification yields three immediate corollaries that explain why Stapp’s work is so important for any theory that links mind and quantum measurement.

1. **Psychophysical identity.** Conscious events and state reductions are ontologically coincident: the reduction is the experiential event.
2. **Participatory agency.** The conscious agent is a “participating observer” who chooses which question to pose (the specification of a projection operator) – and this choice has causal efficacy by selecting which set of potentialities is opened to actualisation.
3. **Attention as stability.** Repetition of the same Process-1 action – what Stapp appeals to via the QZE – stabilises a chosen outcome in the brain, giving a plausible account of sustained attention and intentional control.

Stapp’s formulation is tightly disciplined by the formalism of quantum mechanics and by a methodological modesty: he aims to show how conscious phenomena can coherently inhabit the existing theory without altering its mathematics. That makes his approach powerful, but it also leaves several deep questions unanswered:

1. **The origin problem.** Stapp treats consciousness as a fundamental feature of reality where appropriate physical conditions obtain, but he does not account for how or why the first conscious agent would arise out of mindless matter. The existence of a fundamental psychophysical correlate is taken as a primitive rather than derived.
2. **The metaphysical ground.** There is no explicit account of a metaphysical ground from which both potentiality and actualisation issue. The formalism tells us how to model Process 1 and Process 2, but not why there is processual bifurcation that yields a lived interior at all.
3. **Cosmological scope.** Stapp’s analysis is local and neurophysical: it explains how discrete conscious events relate to brain dynamics. It does not, and was never intended to, provide a cosmology in which the physical world itself is brought into existence through the operation of consciousness.

Orchestrated Objective Reduction (Orch-OR)

The Penrose–Hameroff theory of Orchestrated Objective Reduction (Orch-OR) is the most famous attempt to link quantum physics with consciousness. Like Stapp's model, it takes the quantum measurement problem seriously and treats collapse as an ontologically real process. However, its orientation is more physicalist than participatory, locating the mystery of consciousness not in metaphysical choice but in the structure of spacetime itself.

Roger Penrose introduced the concept of Objective Reduction (OR) to address what he saw as a fundamental incompleteness in quantum mechanics. Superpositions of mass distributions, he argued, create tiny but real distortions of spacetime geometry. When the energy difference between alternatives exceeds a critical gravitational threshold, the superposition becomes unstable and collapses spontaneously. This collapse is not caused by observation but by the self-inconsistency of spacetime curvature. Penrose identified this as an objective physical event – a natural limit to linear evolution. Stuart Hameroff extended this insight into neurobiology, proposing that quantum coherence occurs in the microtubules that scaffold neurons. These cylindrical protein lattices, he argued, can support coherent oscillations shielded from thermal noise. When their collective gravitational self-energy reaches Penrose's threshold, an OR event occurs – and, according to Orch-OR, this event is a discrete moment of conscious experience. Consciousness, in this view, is a sequence of orchestrated reductions within the brain's quantum cytoskeleton, producing a unified stream of awareness.

In one sense, Orch-OR does imply a threshold mechanism: a limit at which quantum potentiality becomes classical actuality through gravitational instability. But this threshold is physical, not metaphysical. It marks the boundary of linear superposition, not the boundary between worldless process and embodied perspective. OR specifies when a collapse must occur, but not when or why such collapses become subjective. The transition from proto-conscious moments to lived interiority remains unexplained.

Another complication is that microtubules themselves are not unique to brains; they exist in virtually all eukaryotic cells. If OR events within microtubules generate consciousness, then every living cell would host some degree of awareness. Hameroff occasionally hints at such a panpsychic implication, but Orch-OR offers no clear account of how isolated OR events combine into a single, self-coherent subject. The model therefore lacks a categorical criterion distinguishing conscious systems from nonconscious ones. Penrose's gravitational threshold can be reinterpreted as a physical analogue of a deeper metaphysical boundary, but it cannot itself serve as the boundary between possibility and perspective.

Several key issues remain open within the Orch-OR framework:

1. **The Binding Problem:** How do countless local OR events cohere into a unified field of experience? The theory invokes microtubule synchrony but never explains phenomenal unity.
2. **The Hard Problem:** Why should collapse *feel like* anything? Objective reduction is a physical process, but Orch-OR does not justify the leap from collapse to qualia.
3. **The Evolutionary Problem:** If microtubules underlie consciousness, why are only certain organisms – those with complex nervous systems – aware? The model lacks a clear evolutionary demarcation.
4. **The Teleological Problem:** OR is spontaneous and gravitationally determined, not guided by value, intention, or purpose. It cannot account for agency or meaning.

The Quantum Convergence Threshold (QCT)

None of the above thresholds can simply be slotted into my cosmological-metaphysical system. When I was writing my last book (*The Real Paths to Ecocivilisation*), the least bad options were QZE and Orch-OR. But for my cosmological model to work, there has to be something very special about animal brains: some sort of property which explains why and how brains are necessary for consciousness, and which also explains what makes conscious animals uniquely involved in wave function collapse. Here is the relevant section, from page 221:

An important question remains unanswered. What exactly is it about animal brains that enables them to embody the Participating Observer when nothing else does? It must be some sort of physical property, and presumably it is relevant to wave function collapse in quantum mechanics, but that's about as much as we can say. Nagel's suggestion that we lack crucial information about the brain certainly applies here (I hold no strong opinion about Penrose and Hameroff's microtubule theory). However, if panpsychism and materialism are both false then it follows that there must be *something* very special about brains – there must be some physical reason why animals alone are conscious, and it can't be anything purely to do with complexity or information processing or any of the other materialistic suggestions.

I can no longer fully agree with the above passage. Firstly it would have been more accurate to ask why *nervous systems* are necessary for consciousness, for jellyfish do not have brains, and at that point I still had no clear answer to the jellyfish question. But the more serious mistake was to equate an informational threshold with physicalism, although at the time I had never considered that this might be wrong. While IIT and GNWT are indeed materialistic – they do not offer a solution to the Hard Problem, and there is nothing metaphysical about them – it does not follow that an informational threshold implies physicalism. This is only true if you start out with an assumption of physicalism (that information has to be encoded within a physical structure), and I was already a neutral monist. If you aren't bothered about the Hard Problem and the Measurement Problem then a physicalist assumption is unproblematic, but if you believe that the fundamental level of reality is itself informational, then why should a purely informational threshold imply physicalism? Retrospectively this is obvious, but it is an insight that eluded me until the thinking of somebody else intervened in my own thought processes.

After RPE was finished, I began searching for places on the internet to talk about it. In one of these places I ran into somebody who was introducing a new quantum interpretation of his own to the world – the “Quantum Convergence Threshold” model (QCT). His name is Gregory Capanda, and he described his theory as a new version of physical collapse, finally able to consign both MWI and CCC to the history books. To me, the situation looked rather different, for regardless of Greg's firm commitment to materialism, QCT actually treats reality as if it is made of information rather than matter. As a theory of wave function collapse it is incomplete, for while it offers an explanation as to when and why the wave function must collapse, it says nothing at all about what collapses into what, or why or how one particular outcome is selected. It merely says that the situation cannot continue as a superposition, so collapse is required. It is therefore an interpretation that itself needs further interpretation. What mattered from my perspective was that threshold it specifies for wavefunction collapse is *specific to information processing systems*, the most obvious and relevant example of which are brains, regardless of the fact that Greg was unclear about what, if anything, his theory had to say about consciousness. This was the beginning of an important development in my own thinking: maybe what we're looking for isn't a physical property at all, even if depends on a physical structure. Maybe what actually matters is the information itself.

I made no attempt to modify QCT. I had enough to deal with planning a book launch at the same time as trying to make progress with my own ideas. QCT came closer to doing the job than

any other theory I'd come across, so I tried to convince its inventor that his theory was inconsistent with materialism and actually *needed* a broader metaphysical model. I was fully aware of the unfinished status of my own theory at that point, but I also wanted to be able to discuss it in public, having kept it under wraps for many years. I was fearful of somebody stealing my core idea and claiming it as their own, so I produced two documents¹¹ which describe an evolving Two-Phase cosmology, using a combination of QZE and QCT as the collapse mechanism. However, when I began working on the book you are currently reading, I had to take a critical look at both QCT and QZE. What I actually required was something which retained the idea of an informational threshold, dispensed with QCT's idiosyncratic mathematics, and also addressed the questions that Stapp's QZE leaves unanswered. QCT had been close enough to the correct puzzle piece for me to make major progress elsewhere in the model (especially cosmology), but it has to be clunkily combined with QZE to make this work. So I asked myself what might actually fit perfectly. What sort of threshold was I actually looking for here?

The idea I needed to retain and develop was that the collapse of the wavefunction is not a passive process but one that occurs when a system reaches a *threshold of commitment to a particular future*. In Greg's terms, this is the Quantum Convergence Threshold – the point at which a coherent superposed system must collapse because it is about to make *an irreversible choice that excludes incompatible futures*. This reframes collapse not as a response to measurement per se, but to a decision about what kind of world will persist into the future. QCT ties collapse to agency: **a system that must decide between real, mutually exclusive futures cannot continue in a superposition**. Collapse is therefore endogenous to a “real choice”, but how is this choice related to consciousness? What are the correct inter-relationships between consciousness, free will and wavefunction collapse? The relationship between consciousness and free will was particularly troublesome, because it seems like these must be two different things: something like passive vs active consciousness. Consciousness seems analogous to an input and free will to an output, and yet it is clear that they are deeply entwined. To make further progress I had to start again, without QZE or QCT, and figure out what a bespoke puzzle piece would look like. This is what I now call the *Embodiment Threshold*.

The philosophical context

Firstly I do need there to actually be a threshold, or the cosmological model I'm describing falls to pieces. The 'embodiment' in question is the point where the Void can become an individual 'view from somewhere,' because a physical structure (a living brain) exists which has whatever is required for embodiment. It needs the structural precondition for the realisation of Brahman as Atman. This structure is what is required to make actualisation logically possible, not to receive some sort of signal from Phase 1.

It is also essential that the model, including the Embodiment Threshold as I am going to define it, is internally coherent and consistent with known science. As far as I am aware it passes both these tests, and if anybody wants to claim otherwise then the burden of proof is on them to demonstrate why the model fails on logical or empirical grounds. As things stand, there is no burden of proof on me to prove the model is true, because it has the same status as all the other interpretations of quantum theory, and all theories of consciousness: it is another metaphysical framework, not yet an empirical scientific theory. There are no other philosophical frameworks which offer a coherent, integrated solution to the Hard Problem and the Measurement Problem, can explain how consciousness evolved, why the constants are fine tuned, and all the other problems this model addresses in a new way. Therefore, if this theory is indeed consistent with both logic and empirical science, and so qualifies as a new philosophical interpretation of quantum theory, and if it

¹¹ *The Participating Observer and the Architecture of Reality* ([10.5281/zenodo.15618749](https://zenodo.org/record/15618749)) and *The Reality Crisis* ([10.5281/zenodo.15823609](https://zenodo.org/record/15823609))

also solves far more outstanding anomalies than anything else, then it has already served a useful purpose. It would be the first empirically adequate theory of everything. Further empirical proof (something fundamentally new rather than resolution of existing anomalies) would be wonderful, but I arrived at this position by following reason, not by doing new science. It is a Kuhnian revolution of the strongest kind, but it is primarily philosophy, not science, and I wish to stress that the philosophy must come first. If my proposed threshold turns out to be incorrect, for whatever reason, the wider theory will still stand, and the search must begin for a new threshold. The only thing which could conclusively falsify it is some kind of proof that no such threshold is possible.

The other element of the mechanism – the question of exactly how the threshold process selects a particular timeline from the physical possibilities, especially if it doesn't just involve individual free will but also involves multiple agents and/or wider-scale phenomena such as synchronicity (which I will be discussing in detail in Chapter 16) – can also be left open, at least to some extent. I am categorically *not* trying to describe a new spiritual system, let alone prove such a thing is true. It is possible that there can be multiple different accounts of how this threshold works and what metaphysical entities are involved – a Buddhist version, a Hindu version, a Taoist version, a New Age version, even an eco-Christian version or a Vita Sapien version. From my perspective, leaving room for different spiritual interpretations in this way is a positive attribute of this system, not a failing. Presuming such a thing actually exists, the threshold itself is a structural feature of reality. The spiritual interpretations concern what, if anything, stands behind that structure. However, I am obliged to provide at least some idea of how this threshold works.

Chapter 10: The Embodiment Threshold

All the previous candidates for the threshold failed because they tried to define consciousness from the outside. I am going to make a case that the real threshold has to be defined from the inside, and I begin with the observation I made in the introduction to this book.

What consciousness does

There may be no non-controversial *materialistic* answer as to what consciousness does, but with a combination of neuroscience and information directly available to us as individual conscious beings, the *function* of consciousness could scarcely be more obvious: consciousness is a process whereby we select our preference among the range of futures we believe to be possible. It begins with the construction of a model of an objective world, within which our bodies exist. This already implies a “self” – an “I”. We see ourselves as a coherent unit – an identity which persists over time. It feels like I'm the same me as I was ten years ago, and in ten years time (assuming I am still alive) I'll still be that same me. Many people have argued that this self is some sort of illusion, but if so it is a very convincing one. It is only because we think of ourselves as persisting entities that we have any motive at all to make mental models of future possible worlds and assign value to the various options. If there is no me then there is nothing to care about the future.

This is another reason why MWI seems crazy: if we actually made every physically possible decision then every time we walked near the top of a cliff, wouldn't there be certain timelines where we spontaneously decide to jump off? Nothing in the laws of physics prevents us from doing so – it is just that we usually choose not to because we place a very low value on the future which would result from that choice.

Perhaps this example is too extreme, but something like this is true of almost everything we do – when children play they are practising how to do this. When we play games we are trying to figure out how to win the game. When we're at work we are figuring out how to continue being paid, and we're being paid to select a particular future, decided by the nature of the work. It is as true for the fine movements required to catch a ball as it is to the decisions I am making when writing the words of a book like this – the future I prefer is one where the ideas I'm trying to communicate are understood by its readers. Even when we are relaxing we are continually making decisions about what is happening around us, and what we want to be doing in ten seconds, ten minutes, or ten hours.

So from our perspective as conscious beings ourselves, the purpose of consciousness appears to be something like this:

- (1) To create and maintain a model of the world we find ourselves embedded in, with ourselves in it as an entity which persists over time.
- (2) To predict the future course of events, which is necessarily a range of possible futures, of differing probability, rather than just one specific future.
- (3) To assign value and meaning to these possibilities, in an attempt to actualise the best possible future, or to avoid the actualisation of the worst ones.

A good demonstration of why this sort of definition is justified is provided by the sea squirt, which consumes its own brain once it has permanently attached itself to a rock. While it is still a tadpole-like free-swimming juvenile, the animal requires a simple brain to navigate and decide among different possible futures: where to swim, what to avoid, where to anchor, etc... Once it has fixed itself in place and its future behaviour can be governed purely by reflexes, that decision-making capacity is no longer needed, so it digests its own brain, which is now more useful as food than it is for making decisions.

The above definition of the purpose of consciousness is fully consistent with all of the available neuroscience, and our direct experience of our own consciousness. Two notable groups of people are likely to object. The first are epiphenomenalists – people who believe consciousness has no causal power and therefore cannot be involved in decision-making – but epiphenomenalism suffers from a fatal problem: if consciousness has no causal effect on brains, then how can brains know anything at all about consciousness? The other common objection cites the Libet experiments, which seem to rule free will out entirely. On this view, the brain makes the decision completely on its own, and consciousness is only notified later. I will have more to say about this in chapter 12, for the results do not actually demonstrate what they appear to demonstrate. These objections notwithstanding, let us assume the purpose of brains and consciousness is as described above.

MWI, minds and the threshold

We do not live in a world in which people randomly jump off cliffs or perform every physically allowed action, and there is no reason to believe our own history is an exception. Yet this is precisely what MWI-style reasoning appears to imply. On that view, whenever a decision is made, reality branches into multiple futures and different versions of the decision-maker proceed in each branch. If taken seriously, this entails that every possible choice is realised somewhere and that the unity of the self is an illusion sustained only by branch-relative ignorance.

Two-Phase Cosmology turns this picture inside out. Rather than being the point at which minds split across diverging timeline, the Embodiment Threshold is the point at which a unified representational subject makes further unitary evolution impossible. Rather than consciousness fragmenting to accommodate a branching reality, reality is forced to stop branching because a singular subject cannot remain coherent across incompatible futures. Wavefunction collapse, consciousness and free will are three different names we have given to this process, in each case without understanding how that process fits into a coherent model of reality. This is not a physical trigger in the usual sense. After more than a century of effort, no experiment has demonstrated a physical variable that initiates collapse. The mechanism described here is openly logical and metaphysical rather than empirical. Collapse occurs because a contradiction arises within the internal organisation of an embodied system, not because a physical threshold is crossed in the wavefunction itself.

The Embodiment Threshold is reached when a biological system inside Embodied Reality develops predictive structures that reference its own possible futures and assign value to them in incompatible ways. Below this point, all representational content can coexist passively within unitary evolution. The system can model multiple possibilities without committing to any of them as its future. Above this point, the system's representational organisation becomes unified enough that it must treat those possibilities as alternatives for a single self. When incompatible valuations are assigned across locally entangled alternatives, the system's own self-model can no longer be extended coherently across them. Continued unitary evolution would require the subject to become internally contradictory. At that moment, unitary evolution cannot continue at that location. This is not because physics breaks down, but because the representational subject cannot. The biological system is forced into actualisation because the self cannot be split. The threshold is therefore a constraint imposed by logic rather than by physical law.

I must emphasise that the Embodiment Threshold applies to embodied systems in Phase 2, not to abstract histories in timeless possibility. Phase 1 contains no evolving brains, no information processing, and no proto-subjects. The alternatives relevant to the threshold are Phase-1-type possibilities instantiated locally within an embodied system's entanglement structure, but the contradiction itself exists only within the system's Phase 2 representational organisation. Collapse does not occur in Phase 1, nor is it prepared or biased there in any way. What crosses the threshold is not a material brain as such, but a particular form of internal organisation within a living brain. Earlier organisms could react reflexively to stimuli, but they lacked a unified self-model capable of representing its own future states as open alternatives. What distinguishes a threshold-crossing system is the emergence of a coherent, indivisible structure of perspective and valuation: a self that is aware that different futures are physically possible for its own body and for the surrounding world, and that can care which of those futures occurs.

This self-structure is informational rather than material, but it is instantiated in biological hardware. It is not a physical object, nor a simple data structure, but an integrated predictive organisation that can span multiple possible futures. In this respect, brains function less like classical machines that follow a single trajectory and more like systems capable of maintaining superposed possibilities until commitment becomes unavoidable. Once such a system becomes aware of mutually exclusive futures and assigns value to them, it cannot continue to exist in superposition.

This is reflected directly in subjective experience. We are constantly aware of multiple possible actions and outcomes. We can deliberate, hope, or wish. What we cannot do is choose contradictory futures at once. One cannot choose to marry Alice and also choose to marry Bob. This is not merely a limitation of introspection; it reflects a structural fact about the self. The unity of the "I" is not something that can branch. The moment a choice becomes real, it excludes its alternatives. We are instinctively and acutely aware of all of this.

This explains why the mind-splitting picture of Many-Worlds feels so deeply wrong. It conflicts with the lived unity of consciousness because it violates the conditions under which a representational self can exist at all. A self that fragmented across incompatible futures would cease to be a self. The Embodiment Threshold explains why we experience one coherent stream of awareness rather than many diverging ones, even as the physical world presents a vast space of possibilities. The unity of experience is preserved not by denying multiplicity, but by recognising that a singular subject cannot inhabit it.

Formal statement of the Minimum Conditions for Conscious Perspective

Let an agent be any **physically instantiated system within Embodied Reality (\mathcal{R})**. An agent possesses a *conscious perspective* (that is, there is something it is like to be that agent) if and only if the following conditions are jointly satisfied.

1. Unified Perspective.

The agent maintains a single, indivisible internal model of the world that includes itself as a coherent point of view persisting through time. This representational unity is not merely functional integration but logical indivisibility: it cannot be partitioned into incompatible submodels without the loss of subjecthood. A system whose self-model could bifurcate into independent continuations would not constitute a conscious perspective.

2. World Coherence.

The agent's internal model is in functional coherence with at least one real physical state of the external world. This coherence may be local, as in the state of the agent's own body and immediate environment, or extend to broader regions of reality. Valuations directed toward physically unrealisable states do not qualify (or are less effective), since agency requires that at least some evaluated futures be genuinely possible.

3. Evaluation.

The agent can assign value to possible future states of itself or the world, enabling comparison between alternatives. These valuations are not mere reward signals or reactive preferences but reflect the agent's own perspective on what ought to occur. Without valuation, there is no meaningful distinction between futures and therefore no basis for choice or commitment.

4. Non-Computability.

At least some valuations are **non-computable in the Turing sense**. Following Penrose's argument, conscious judgment cannot be exhaustively captured by algorithmic computation. These non-computable evaluations introduce genuine openness into the agent's decision process and prevent the reduction of conscious choice to deterministic or stochastic rule execution. Without non-computability, the agent's behavior would be fully fixed by prior physical states, and the appearance of choice would be illusory.

These four conditions specify the minimal structural and functional requirements for a conscious perspective within Phase 2. When they are jointly satisfied in an embodied system that is locally entangled with multiple physically possible futures, a new constraint arises. The agent's unified self-model must treat those futures as alternatives for a single subject, while its non-computable valuations may rank them incompatibly. Continued unitary evolution would therefore require the agent to sustain mutually contradictory commitments across those alternatives. At this point, unitary evolution cannot be maintained at the level of the agent's representational organisation. Collapse occurs as the logical resolution of this incompatibility, selecting a single embodied continuation that preserves the agent's unified perspective. This process is not triggered in Phase 1, nor is it prepared or biased there in any way. Phase 1 contains only timeless possibilities; the contradiction and its resolution arise entirely within Embodied Reality.

In summary, consciousness, free will, and collapse are not separate phenomena. They are three complementary descriptions of a single Phase 2 process: the resolution of representational contradiction in an embodied system whose unified, non-computable valuations cannot be coherently extended across multiple incompatible futures. A system that satisfies conditions (1)–(4) is collapse-competent; a system that fails any one of them is not.

The storm metaphor

Ordinary consciousness feels continuous: a steady flow in which perceptions, thoughts, and intentions seem to unfold seamlessly through time. Yet, upon closer examination, this continuity is not a property of physical time itself, but of the way consciousness integrates a sequence of discrete moments into a coherent whole. Philosophers and psychologists have long called this interval the **specious present**: the brief window of lived time in which successive events are experienced as part of a single, unified now. The specious present is a field of simultaneity stretched across a short duration – no more than a few seconds, probably much less. Within this window, multiple events coexist in immediate awareness: the note that is fading, the note that is beginning, and the expectation of the next note in a melody. Conscious experience is a dynamic synthesis. The living present is a wavefront of integration.

In 2PC, this living present corresponds to what may be called a storm of micro-collapses: a cascade of local actualisations through which portions of the background superposition are momentarily stabilised. Each micro-collapse is a tiny act of realisation, in which one physically possible configuration becomes actual where representational coherence requires it. These events form a coherent pattern, overlapping and interacting across the specious present, much as eddies and vortices compose the moving texture of a storm.

The storm metaphor captures the essential features of lived consciousness. It is dynamic rather than static; it is a process rather than a thing; it exhibits local coherence amid global

indeterminacy, with stability arising only through continual renewal; and it integrates activity across multiple scales. Countless micro-events combine to form the fluid unity of subjective experience. In this view, the specious present is not merely a psychological artifact but an ontological structure: the timescale over which local aspects of reality are continually stabilised through ongoing collapse activity. The “flow” of consciousness is the temporal cross-section of this activity – the living wave of embodiment where potential becomes actual. Each moment of awareness is a pulse of value-realisation sustained by the Void’s ontological grounding. Continuity of consciousness, then, is not the persistence of a substance but the persistence of a pattern.

Attention and will are the shaping winds of this storm. They do not conjure collapses out of nothing, nor do they directly select outcomes. Instead, over the span of the specious present, they modulate the pattern, timing, and weighting of micro-collapses, biasing the trajectory of experience toward some continuations rather than others.

The storm metaphor also clarifies the limits of consciousness. A storm can stretch across vast regions if conditions align, but it cannot cover the entire globe. Likewise, collapse processes can extend across entangled systems, but they cannot span the whole cosmos, as occurred only at the original transition when historical Phase 1 gave way to historical Phase 2. Coherence gives the storm its reach; decoherence disperses it into background noise. And like every storm, an individual consciousness eventually dissipates. When the biological conditions that sustain the pattern break down, the storm ceases to be. The self (the weather system) and the soul (the Void’s grounding of that system) disappear together. Nothing persists but the wider field of possibility, out of which new storms may someday form.

Self-Models and the Unity of Subjectivity

A growing body of work in philosophy of mind and cognitive science converges on the idea that subjectivity does not arise from a metaphysical essence but from a particular kind of representational organisation. On this view, the self is not a substance but a process: a structured, embodied model through which a system represents the world and itself within it.

Thomas Metzinger’s Self-Model Theory of Subjectivity provides one of the most influential formulations of this idea. Metzinger argues that there is, strictly speaking, “no self,” only a transparent phenomenal self-model (PSM) generated by the brain. This model integrates bodily state, perspective, intention, and agency into a unified point of view. Because the model is transparent (its representational nature is not itself represented) it is not experienced as a model at all. It is simply lived as “me, here, now.” This transparency explains the phenomenal unity of subjectivity: the system cannot access or manipulate the self-model as an object, and therefore cannot experience itself as divided into multiple simultaneous selves.

A complementary account emerges from Antonio Damasio’s work in neuroscience. Damasio distinguishes between the proto-self of basic bodily regulation, the core self of moment-to-moment conscious presence, and the autobiographical self constructed through memory and narrative. The core self, in particular, is a transient, embodied process that arises from the integrated regulation of the living organism. Although dynamically constructed and fragile, it functions as a single centre of experience. Its unity is not imposed from outside but arises from the organism’s integrated biological organisation.

Phenomenological approaches reach a closely related conclusion. Shaun Gallagher distinguishes the minimal self (the immediate, embodied “I” that anchors experience) from the narrative self extended across memory and social context. Dan Zahavi likewise emphasises the intrinsic “mineness” (ipseity) of experience: every conscious episode is given as belonging to a single subject. This first-personal character is not an added feature but a structural property of experience itself. For these thinkers, the minimal self is not a theoretical posit but a phenomenological invariant: experience is always presented from one perspective and cannot be shared or jointly occupied.

An even broader framework is provided by the enactive approach developed by Francisco Varela, Evan Thompson, and Eleanor Rosch. On this view, cognition is not the passive construction of internal representations but the active enactment of a meaningful world through embodied interaction. Subjectivity arises from the dynamic coupling of organism and environment. Thompson has further argued that this participatory structure grounds the continuity of the self: conscious perspective is inherently situated, embodied, and world-involving, not detachable or freely duplicable.

Taken together, these lines of research converge on a single structural insight. Once a system achieves a transparent, embodied first-person perspective (a phenomenal self-model, a core self, or a minimal “I”) its subjectivity is unified in a way that cannot be partitioned without ceasing to exist as a subject. The unity described by Metzinger as transparency, by Damasio as the immediacy of the core self, by Gallagher and Zahavi as the indivisibility of first-person presence, and by Varela and Thompson as participatory coupling, all point to the same constraint: first-person experience is singular.

2PC extends this convergence beyond phenomenology and cognitive architecture. The Embodiment Inconsistency Theorem states that the indivisibility of subjectivity is not merely a feature of experience but a metaphysical constraint on how reality can evolve once such a subject exists. In an embodied system within Phase 2, a unified self-model cannot be coherently extended across multiple incompatible futures. When representational unity and valuation meet unresolved physical alternatives, unitary evolution becomes untenable. Collapse is therefore required, not as an added mechanism, but as the necessary resolution of a contradiction introduced by the existence of a singular, first-person perspective.

In this way, contemporary theories of self-models do not merely describe how the self appears. They implicitly delineate the conditions under which a self can exist at all. 2PC makes this implication explicit: the unity of subjectivity is not only lived, but enforced by the structure of reality itself.

Preliminaries: From Possibility to Embodiment

In 2PC, reality has two ontological regimes. Phase 1 is *Timeless Possibility*: the complete set of physically consistent cosmic histories, denoted Ω . Nothing evolves inside Ω , and nothing within it leans toward embodiment. It is a static field of alternatives, containing multiplicity without subjectivity, valuation, or commitment. Phase 2 is Embodied Reality, denoted \mathcal{R} . \mathcal{R} is the unique instantiated history in which events occur. Everything that happens (consciousness, valuation, agency, and the appearance of collapse) occurs inside \mathcal{R} . Phase 2 arises when a conscious subject exists whose internal valuations cannot remain jointly coherent across the physically possible alternatives with which it is locally entangled. At that point, embodiment is not chosen or caused, but becomes ontologically required.

The Embodiment Threshold (ET)

The Embodiment Threshold is the point at which a system gives rise to an internally unified representational structure whose continued coherence requires the resolution of incompatible valuations across locally entangled, decoherent alternatives. At ET, a single prospective subject is forced to assign mutually incompatible values to alternative future continuations that cannot be jointly realised by that same subject. As a result, the system’s self-model can no longer be coherently extended across all of the branching possibilities it simultaneously references, making ontological collapse necessary.

Formally, let a system S be characterised at a pre-symbolic level by an internal informational organisation, written as $IS(t)$. This organisation is not a physical state description and does not yet

involve belief or semantic content. Instead, it consists of mesoscopic predictive structures capable of supporting valuation under the assumption of a single continuing subject.

Let VS be a valuation map that assigns to each such predictive structure a value in an abstract valuation space V . In simple terms, VS maps elements of $IS(t)$ to values in V .

The Embodiment Threshold is reached when these valuations can no longer be jointly satisfied under the constraint that they all belong to one unified subject. To make this precise, consider a decoherent set of locally entangled alternative continuations, written as $\{\alpha_i\}$, to which the same representational subject assigns valuations. ET occurs when there is no coherent extension of the self-model that allows all of these assigned valuations to be realised by a single subject. In this case, the joint satisfiability condition fails.

In shorthand, ET occurs when:

$Joint_Satisfiability(\{\alpha_i\}, VS) = FALSE$

This failure is not due to physical incompatibility between the alternatives themselves. Instead, it arises from a referential contradiction: the representational "I" cannot be coherently duplicated across incompatible valuation contexts. Ontological collapse is therefore forced at or before ET as a condition of representational coherence, thereby instantiating Phase 2 embodied reality.

The Embodiment Inconsistency Theorem (EIT)

The Embodiment Inconsistency Theorem explains why embodiment must occur once a system reaches the Embodiment Threshold. It is a meta-theoretic consistency result rather than a dynamical law. It does not rely on measurement, physical collapse mechanisms, environmental decoherence, or any modification of quantum dynamics. Instead, it follows from a logical constraint: a single referential "I" cannot coherently sustain stable valuations across mutually incompatible futures.

Theorem:

Let a system S satisfy the following axioms:

- **Valuation:** The system assigns intrinsic valuations $VS(\mathbf{hi})$ to physically possible future continuations. This is defined by a valuation map $VS : \{\mathbf{hi}\} \rightarrow V$, where each hi corresponds to a decoherent alternative locally entangled with S and describable as a Phase 1 admissible continuation
- **Entanglement:** The system's informational degrees of freedom are distributed across multiple decoherent alternatives. Its future self-model spans a family $\{\mathbf{hi}\}$ with stable decoherence relations.
- **No-Overdetermination:** The universal wavefunction provides no rule by which a single subject can reconcile incompatible intrinsic valuations across all \mathbf{hi} . No physical law selects a unified continuation for a valuative subject in advance.
- **Ontological Coherence:** A single subject cannot inhabit futures that assign incompatible intrinsic valuations to its own continuation. This is governed by a coherence predicate $C(\{VS(\mathbf{hi})\})$. If the valuations across the set of histories are mutually inconsistent for a single referent, the predicate returns **FALSE**.

Conclusion

The Embodiment Threshold is reached at the exact moment the coherence predicate fails:
 $C(\{VS(\mathbf{hi})\}) = FALSE$.

At this point, no globally consistent assignment of a single subject across all \mathbf{hi} is possible. The

contradiction lies in the system's own representational dynamics, not in the universal wavefunction. Because reality cannot host a single referent with mutually inconsistent successor states, embodiment becomes necessary. One history is instantiated for **S**, and the system becomes embodied within that history. EIT shows that the existence of a unified subject is incompatible with the continued unitary evolution of its entangled alternatives once valuations diverge beyond the threshold of coherence.

Relationship to Bell's Theorem and Conway-Kochen's Free Will Theorem

The Embodiment Inconsistency Theorem can be understood as a metaphysical analogue of well-known quantum impossibility results. Bell's theorem demonstrates that no local hidden-variable theory can reproduce quantum correlations, and Free Will Theorem shows that no fully deterministic mechanism can constrain the choices of certain systems without conflicting with the structure of quantum theory. The EIT extends this style of reasoning to self-referential, value-bearing systems: no such system can remain uncollapsed while maintaining unified valuations across incompatible alternatives without generating internal logical incoherence. In this sense, consciousness is anchored in necessity: the very structure of a unified, valuative subject enforces the conditions for collapse.

Formal Characterisation of the Embodiment Threshold

Let a Phase 2 system **S** be locally entangled with a family of decoherent alternatives, written as **{hi}**. Each alternative carries a decoherence weight **pi** derived from the diagonal of the decoherence functional. These weights describe the relative presence of alternatives within the superposition but do not cause collapse; they define the domain over which the system's self-model must maintain representational coherence.

The internal organisation **IS(t)** provides valuations **VS(hi)** in an abstract valuation space. Representational coherence requires that there exist a single valuation vector **v*** that could serve as a unified perspective across all relevant alternatives **hi**.

We define the inconsistency index **Lambda_S** over the set **{hi}** as follows:

$$\Lambda_S(\{h_i\}) = \min_v [\sum_i (p_i * \text{distance}(V_S(h_i), v)^2)]$$

$$\Lambda_S(\{h_i\}) = \min_{v \in \mathcal{V}} \sum_i p_i \|V_S(h_i) - v\|^2$$

This quantity is a mathematical summary of the degree to which the system's valuations fail to admit a single coherent perspective. The minimising vector **v*** represents the best conceptual compromise the system could maintain across the alternatives. The Embodiment Threshold, denoted **t***, is reached when **Lambda_S** exceeds the system's tolerance for representational divergence, **Theta_S**. This tolerance is intrinsic to the architecture of the system's self-model and defines the maximum divergence compatible with a unified subject. When **Lambda_S > Theta_S**, representational coherence fails. A single subject cannot be coherently extended across the family **{hi}**. At this point, a single embodied continuation becomes metaphysically necessary, and one Phase 2 history is realised for the system. The contradiction arises entirely within the Phase 2 system's own predictive organisation.

What the Embodiment Threshold Is *Not*

The Embodiment Threshold is not a microtubule resonance, an integrated information measure, a global broadcast, or a gravitational instability. It is the precise moment when a physical system

becomes the kind of thing that cannot be multiplied across decoherent alternatives without internal contradiction. The threshold resides entirely in the system's representational self-model, not in any additional quantum feature.

This framing clarifies the core architecture of 2PC:

- Consciousness is the resolution of representational inconsistency.
- The Void is the ontological ground, not a chooser: it makes the resolution real, but does not actively select outcomes.
- The Embodiment Threshold provides the model with a metaphysical anchor by specifying when collapse is required to preserve ontological coherence. What happens *within* that collapse (namely, which specific future is instantiated) can be treated separately, leaving room for psychological, spiritual, or empirical interpretations.

Relationship to Orch OR and Stapp

Penrose and Stapp both identified an important structural coincidence: moments of conscious experience appear to align with physical state reduction. Penrose sought this alignment in objective gravitational thresholds; Stapp located it in projection events associated with attention. Both approaches helped legitimise the idea that indeterminacy and experience are linked. What neither account explains is why a collapse should constitute a point of view at all. They describe when reduction occurs, but not how reduction becomes subjectivity. In 2PC, collapse becomes necessary only when a predictive system reaches the Embodiment Threshold – when incompatible valuations can no longer coexist within a single representational “I.” Unitary evolution fails for internal, ontological reasons, not because a physical parameter is exceeded. Reduction is therefore not merely a physical event but the instantiation of a coherent perspective.

The Hard Problem (#22) and the Even Harder Problem (#23)

The Hard Problem of Consciousness is a problem only under materialist or physicalist assumptions. From the perspective of 2PC it is readily resolved: consciousness is not produced by matter per se but arises from the instantiation of a unified self-model that grounds Phase 2 embodiment.

The Even Harder Problem, by contrast, concerns the lack of a framework capable of uniting competing metaphysical perspectives. This problem remains open: I cannot claim to have solved it until such time as 2PC has convinced a sufficiently broad community that it provides both a coherent foundation for scientific explanation and a complete, internally consistent model of reality.

The General Anaesthetic Mechanism Problem (#24)

A 2PC-style approach to understanding general anaesthesia does not seek a specific pharmacological “switch” that turns consciousness off. Rather, it asks: **what process must fail** for consciousness to cease?

In 2PC terms, general anaesthetics prevent the brain from reaching the Embodiment Threshold. This reframes the empirical question: the blockade is not of awareness in general, sensory pathways, or cortical electrical rhythms. It is of the formation of the mesoscopic pattern of the self-model (the pattern that allows a system to become the referent of its own predictions and to hold valuations coherently across the specious present).

This reframing has practical implications. Instead of searching for a single molecular target shared across xenon, propofol, and ketamine, the focus shifts to what these substances disrupt at the mid-scale architecture of the brain. They interfere with the stability of the representational storm that would otherwise instantiate a coherent subject. Perception, memory, and reflexive processing

remain largely intact, but the brain cannot organise itself into a unified, valuative self. Seen this way, anaesthetics do not reveal what matter produces consciousness; they reveal what a brain must do to support consciousness. They block the very process that LUCAS's system once achieved: the self-referential integration of predictive templates across entangled alternatives.

Chapter 11: Psycheogenesis and the Psychetelic Principle

People have long wondered about abiogenesis – the emergence of life from non-life. In the Two-Phase cosmology, what matters is the first appearance of conscious organisms, and abiogenesis gets thrown in for free, because it is necessarily on the path from non-life to consciousness. I use the term “psycheogenesis” to refer not only to the emergence of conscious life from unconscious life, but to the entire process that made this possible. The resulting explanations are related to the Anthropic Principle – the idea that it is only in cosmoses or timelines where intelligent beings like humans exist that anybody asks questions about how their world came into being. The Psychetelic Principle (“psyche” + “telos”) says that a cosmos is only realised if it can host a subject, and that its history was always destined toward the conditions that let consciousness arise. Psycheogenesis is the long, directional drift of matter toward structures that can carry a point of view, and once that point of view appears, the cosmos locks into the form that consciousness can inhabit. In other words, the Psychetelic Principle provides a mechanism, not just an unsatisfying explanation.

The Cambrian Explosion and the Evolution of Consciousness (#27, #28)

In *Mind and Cosmos*, Thomas Nagel argued that there can be no materialistic neo-Darwinian explanation for how consciousness evolved, and that the only credible naturalistic alternative involves teleology: consciousness was destined to evolve. The Two-Phase Cosmology proposes that the explanation for this apparent teleology is not the undiscovered teleological laws that Nagel suggests we should be searching for, but the same thing that resolves the Hard Problem and the Measurement Problem. The Cambrian Explosion is now permitted to have the explanation that should always have been obvious. The unifying characteristic of the Cambrian fauna was that they actually look like animals as we intuitively understand that word to mean. This is in contrast to very simple animals like sponges, which many people assume to belong to some other branch of life entirely (as did early naturalists). The same applies to the Ediacaran fauna which dominated the world in which LUCAS lived – they were either sessile (they didn't move at all), or their mobility was akin to that of jellyfish.

The Embodiment Threshold clarifies the jellyfish question. They lack brains, and it is hard to see how they can be intuitively aware of different possible futures about which they can make decisions. Their movements are purely reflexive. They are likely to be examples of the most advanced level of behavioural complexity that unconscious animals are capable of. I obviously cannot prove this, but my tentative conclusion is that their nervous system does not cross the Embodiment Threshold. A “central processing unit” is required: a *brain*.

The immediate ancestors of LUCAS must have been stymied by the Frame Problem. Their increasing cognitive complexity would have amplified the combinatorial explosion of potential futures they could model, but at that point they had no mechanism to assign meaning or value. Without the capacity for valuation that only comes with consciousness, effective decision-making would have been difficult. However, so long as reality remained superposed, meaningful choice was metaphysically impossible anyway – all physically possible outcomes existed in parallel, just as believers in MWI think they still are.

The only way for the informational structure to continue coherently in this situation is through a selective mechanism – a mechanism experienced as free choice. This moment is

psychegenesis and the beginning of phase 2. Cosmic history before this point was actualised retrocausally: the entire evolutionary trajectory leading to LUCAS was selected as a unified classical timeline. After this point, evolution could proceed with genuine decision-making capacity. A new sort of phase 1 information structure began to exist. The original phase 1 history, selected for realisation at the moment LUCAS becomes conscious, was a timeless and unchanging block. Now there was a dynamic structure in which the Void was embedded and animals were conscious.

My suggestion for the best candidate we're aware of for LUCAS is a worm-like organism smaller than a grain of rice. It lived at the end of the late Ediacaran Period, around 560 to 555 million years ago, not long before the Cambrian Explosion. It is the earliest confirmed organism exhibiting bilateral symmetry, a through-gut, and a clear front-back axis. It is the simplest known organism that might have been capable of making choices in a metaphysically meaningful way. *Ikaria wariootia* was small, blind, and simple, but for a brief period of time the Earth was its oyster; at the end of the Ediacaran nothing could challenge it. It appeared when the only other complex “animals” were so unlike anything living today that scientists are uneasy using the word “fauna” to refer to them. Maybe they weren't animals at all, but some other branch of life which died out, because it was unable to compete when the Cambrian got going. Under 2PC we can shed some more light on this.

The Ediacaran ecosystem existed at the start of phase 2 of cosmological history, but its origin was in phase 1, and like everything else in phase 1 it was retroactively selected because it was on the path that led to LUCAS. The Ediacaran fauna were the supporting cast of actors which were necessary for there to be a functioning ecosystem for LUCAS to evolve within. On the direct line to LUCAS (which was necessarily one very specific lineage) this evolution was directed by teleology rather than natural selection, but presumably some kind of natural selection did apply to everything *not* on that direct line. Some of these organisms must have had nervous systems but no brains, like jellyfish and comb jellies today. *Ikaria* was the first bilaterian, and with this new symmetry came the capacity to move forward with purpose, to push through the microbial mats of the sea floor in search of nourishment. It was the first organism which had a solution to the Frame Problem, and this must have provided an overwhelming advantage. So *Ikaria* multiplied and spread, and feasted on whatever it liked, since most of the species with which it shared an ecosystem were defenceless against an animal which could think.

This sort of abundance never lasts. Boom leads to bust. Nothing was stopping *Ikaria* from multiplying until the easy food was gone, the microbial mats stripped bare and the sediment was fouled by overpopulation. And then, inevitably, came die-off: dead *Ikaria*, littering the seabed. Any that were able to scavenge the remains survived longest, and a new opportunity was born. The first to discover a carcass had the best chance of feeding and surviving, and in this game, the slow, the hesitant, and the dull-witted perished. To find food in a world where it appeared and vanished unpredictably, some *Ikaria* needed to process more than simple chemical gradients.

Evolution responded, and this time it was natural selection as we understand it: nervous systems began to centralise and specialise. Some neural circuits became more attuned to the predictable – to recognising patterns that indicated a meal or a mate. Others stayed watchful, scanning for novelty – for the unexpected signs of opportunity or threat. The left and right sides of the brain began to diverge, one side categorising and predicting, the other staying vigilant, wary of being trapped by its own expectations. Only essential information was communicated between them, so that each could optimise its own function (see the work of Iain McGilchrist).

What began as scavenging eventually became hunting. While some animals grew better at locating dead flesh, others began to chase the living. Predators emerged, sharpening their senses, growing swifter and more deliberate. The seas, which had been silent and slow, became places of pursuit and evasion. Prey species responded with speed, armour, camouflage or cunning. The arms race between predator and prey accelerated, and with it, the relentless push toward neural complexity. Among the early predators were creatures that would give rise to molluscs and

arthropods, animals with shells, eyes, and limbs. In their brains, the hemispheric divide deepened. One side projected the path of fleeing prey, calculated the angle of pursuit, anticipated where the target *would be*. The other side watched for surprises, for sudden changes in the terrain, or for signs of an ambush in the shadows. The brain became a balance between the known and the unknown, prediction and vigilance. At the peak of this evolutionary escalation stood *Anomalocaris*. It was unlike anything before it: a metre-long predator with powerful swimming flaps, stalked eyes with compound lenses, and great grasping limbs for snatching prey. *Anomalocaris* was the full flowering of attention and decision. Its eyes saw movement and subtle shifts. Its brain processed that input, split between prediction and watchfulness, continually solving the Frame Problem on the fly, deciding in an instant whether to strike, pursue, or veer away.

It took about 40 million years to get from the simple burrower *Ikaria* to the formidable *Anomalocaris*. The story of conscious animal life was, and remains, all about meaningful choice. No amount of speed or strength can substitute for the ability to make the right decisions at the right time.

The Frame Problem (#26) and the Binding Problem (#25)

The Frame Problem is the challenge of how a system knows which aspects of the world are relevant to a situation or action, without having to evaluate *everything*. The Binding Problem is how disparate sensory inputs and neural processes unify into the seamless experience of a single self. Under the model being described neither of these problems seem problematic any more. The introduction of both a structural-informational and ontological subject (the self and the Void) provides a direct answer to both questions – it is no longer a mystery why conscious beings do not suffer from the Frame Problem or the Binding Problem, because this “self” is what binds consciousness together and provides a coherent frame. Consciousness does not exist at all without it. It is not a fragmented set of qualia loosely bound together by physical processes. It is a singular, indivisible experience arising from the Void’s participation in a superposed brain at the moment of wave function collapse. There is only one "I," because for each conscious being there is only one Atman. This model makes clear why humans and animals are so remarkably adept at managing the Frame Problem compared to artificial systems. Unlike mindless machines, they engage in participatory collapse events. This is a fundamentally non-computational process, irreducible to algorithms or data processing alone. It is the metaphysical intervention of the Void that grants animals the ability to make meaningful, holistic decisions under uncertainty.

The Frame Problem is usually framed as a worry about how any system can know what matters in a situation without checking every possible variable, and the binding problem is usually framed as a question about how scattered neural processes come together as one experience. In the picture we are building these issues stop looking like puzzles because the subject is not an after-the-fact construct. The structural subject and the ontological subject arise together, so the field of experience is already unified before any cognitive operation takes place. The storm of micro collapses across the specious present does not stitch fragments together from the outside. It stabilises a single point of view that already carries coherence because the Void grounds every collapse that keeps the self intact.

The Frame Problem falls away once relevance is not computed but lived. A conscious agent does not sort through an implicit list of possibilities. It resolves representational conflict by collapsing only the region of the superposed brain that it is already entangled with, and that region is shaped by its history of valuations, its predictive templates, its attentional habits, and the internal structure of its self model. Relevance flows from the pattern of collapses that sustain the subject rather than from a rule set.

The Binding Problem falls away once unity is not a neural achievement but an ontological requirement. There is one viewpoint because collapse cannot support a split referent. The representational I cannot branch into multiple inconsistent centers of valuation, so experience

arrives as a single field. Sensory features come together because the collapse process already has to respect the coherent subject that the Void stabilises, and the brain's mid scale predictive structures settle into that unity whenever they resolve a conflict.

This is why conscious animals handle context, salience, and meaning with an ease that artificial systems cannot imitate. Their decisions are shaped by participatory collapse rather than computational search, and they do not need an algorithm for relevance or unity because their existence as subjects already gives them both.

The uniqueness of the original psychogenesis event (the collapse of the primordial wavefunction).

The first phase shift event – the one involving LUCAS – is unique in two ways. Firstly it selects an entire cosmic history stretching back to the big bang, not just a small part of it. Secondly it is the original selection, and this cannot be “invoking” or “compelling” the Void to involve itself in the situation. The problem here is that Pythagorean Ensemble of Phase 1 must contain infinite examples of possible structures which arrive at the ET. If each of them had the power to *compel* the Void then it would need to respond to all of them simultaneously (and possibly even repeatedly). This situation would be even more ontologically bloated than MWI, but is easily avoided by assuming that only one cosmos is realised at any one time. If this model is correct then there are an infinite number of realities-in-waiting, and eternity for it to be their turn. I call them “cosmic eggs”. Somehow a particular reality must be chosen for realisation when the previous cosmos has run its course. The determination of which of these cosmic eggs are selected into reality, and in which order, is an interesting question. I am skeptical that it can be random, because I can't imagine where randomness could come from. It is possible that some informational rule determines this original selection (and I will have more to say about this in Chapter 14). Perhaps there is some other explanation. All we can say that it does indeed happen, and our own cosmos was the most recent cosmic egg to hatch.

However, that there is one cosmos/reality does not imply that there can only be one conscious organism, and once a cosmos and timeline has been realised at the collapse of the primordial wave function then any future instances of ET-crossing organisms can indeed “invoke” the Void. If they satisfy the physical conditions, then they will become conscious. This must be true under this model, for it is the reason why any physical organism capable of becoming conscious actually does become conscious. Once Phase 2 has begun then Brahman is always ready to become Atman when a *view from somewhere* is available.

The Psychelic Principle

From a conventional scientific viewpoint, consciousness is an inexplicable enigma. From a Darwinian perspective grounded in survival and reproductive advantage, the emergence of subjective experience is profoundly puzzling. Why should an undirected, mechanical process give rise to inner life, rather than simply more efficient stimulus-response mechanisms? In *Mind and Cosmos*, Thomas Nagel argued that natural selection, as currently understood, cannot explain the emergence of conscious subjectivity, and proposed we search for teleological laws of nature: goal-oriented principles embedded in the fabric of the cosmos. The new model resolves the problems on all fronts. Consciousness is not selected, but ontologically prior to evolutionary competition within our observed universe. Where Nagel suggests that evolution must be guided by teleological laws aimed at producing minds, 2PC posits a structural inevitability: in a potentially infinite quantum cosmos, some branch will arrive at the ET, and there it will remain in timeless incompleteness until such time as the Void resolves the situation by realising that structure, allowing consciousness to emerge and start collapsing possibility into actuality. If and when that happens then that branch becomes a realised cosmos inhabited by conscious observers. This retains Nagel's key insight that

mind is not epiphenomenal or accidental, but dispenses with his hypothetical teleological laws. The apparent directedness of evolution toward complexity and consciousness is a selection effect caused by the fact of consciousness itself being the criterion for observable history. The universe we observe is the one rendered actual by being observed, regardless of probability. This results in something very similar to the anthropic principle, except instead of just saying “If humans hadn't evolved then we wouldn't be here to ask the question”, we're actually explaining why conscious organisms were guaranteed to win the cosmic lottery in this way. I call this “the Psychetelic Principle”.

Why did psychogenesis happen on Earth, rather than somewhere else? The PP tells us that we should expect the Earth to be special, but it doesn't tell us exactly what *is* special about Earth. It makes an empirical prediction: if the model is correct, then there should have been multiple exceptionally improbable events in Earth's phase 1 history. These, if they exist, would be signatures of psychogenesis. There are many examples, the most extreme of which are these four:

1. Eukaryogenesis: The Singular Emergence of Complex Cellular Life

Eukaryogenesis is the origin of the eukaryotic cell via the endosymbiotic incorporation of an alpha-proteobacterium (the precursor to mitochondria) into an archaeal host, and it appears to have happened only once in Earth's entire 4-billion-year history. Without it, complex multicellularity (and thus animals, cognition, and consciousness) would not have emerged. The energetic advantage conferred by mitochondria enabled the explosion of genomic and structural complexity. No similar event is known to have occurred elsewhere in the microbial biosphere, despite vast diversity and timescales. If eukaryogenesis is a statistical outlier with a probability on the order of 1 in 10^9 or worse, it becomes a cardinal signpost of the unique psychogenetic branch.¹²

2. Theia Impact: Formation of the Earth–Moon System

The early collision between Earth and the hypothesised planet Theia yielded two improbable outcomes at once: a large stabilising moon and a metal-rich Earth. The angular momentum and energy transfer needed to both eject enough debris to form the Moon and leave the Earth intact was extremely finely tuned. This event likely stabilised Earth's axial tilt (permitting climate stability), generated long-term tidal dynamics (affecting early life cycles), and drove the internal differentiation which fuels the magnetic field and active tectonics. It's estimated to be a rare outcome among rocky planets – perhaps 1 in 10^7 – and essential for the continuity of biological evolution.¹³

3. Grand Tack: A Rare Planetary Migration Pattern

Early in solar system formation, Jupiter is thought to have migrated inward toward the Sun and then reversed course (“tacked”) due to resonance with Saturn. This migration swept away much of the early inner solar debris, reducing the intensity of late bombardment and allowing small rocky planets like Earth to survive. Crucially, it also delivered volatiles (including water) to the inner system. This highly specific orbital choreography is rarely reproduced in planetary formation simulations. Most exoplanetary systems dominated by gas giants do not preserve stable, water-bearing inner worlds. The odds against such a migration path are estimated to be very high. Some

12 Lane, N., & Martin, W. F. (2010). The energetics of genome complexity. *Nature*, 467(7318), 929–934. <https://doi.org/10.1038/nature09486>

13 Canup, R. M. (2004). Simulations of a late lunar-forming impact. *Icarus*, 168(2), 433–456.
Laskar, J., Joutel, F., & Robutel, P. (1993). Stabilization of the Earth's obliquity by the Moon. *Nature*, 361(6413), 615–617 and Elser, S., et al. (2011). How common are Earth–Moon planetary systems? *Icarus*, 214(2), 357–365, and Stevenson, D. J. (2003). Planetary magnetic fields. *Earth and Planetary Science Letters*, 208(1–2), 1–11.

simulations suggest well under 1 in 10^6 .¹⁴

4. LUCA's Biochemical Configuration

The Last Universal Common Ancestor (LUCA) did not merely represent the first replicator, but a highly specific and robust configuration of metabolism, information storage, and error correction. It was already using a universal genetic code, RNA–protein translation, lipid membranes, and a suite of complex enzymes. LUCA's molecular architecture was a kind of “narrow gate” through which life could pass toward evolvability. Given the astronomical space of chemically plausible alternatives, LUCA's setup may reflect a deeply contingent and rare outcome.¹⁵

Conclusion: Compound Cosmic Improbability as Psyche-genetic Marker

Each of these four events is, in itself, vanishingly unlikely. But more importantly, they are compounded. The joint probability of a single planet experiencing all four – along the same evolutionary trajectory – does indeed render the Earth's phase 1 history cosmically unique. Under 2PC these improbabilities indicate the statistical imprint of consciousness retro-selecting a pathway through possibility space – making a phase transition from indefinite potentiality to a single, chosen actuality.

The Preferred Basis Problem (#20)

In QM, the wavefunction does not come pre-packaged as a world of tables, bodies, and trajectories. Mathematically, it can be expressed in indefinitely many equivalent bases, each corresponding to a different way of carving up physical possibilities. Nothing in the formalism tells us why experience should correspond to this particular decomposition (objects at definite locations evolving over time) rather than some radically different one. This is the Preferred Basis Problem: why does experience appear in one specific structural frame when the underlying theory permits infinitely many?

Decoherence is often said to solve this problem by showing how interactions with the environment suppress interference between certain states, stabilising quasi-classical patterns. Decoherence explains why some structures behave as if they were classical, and why superpositions become dynamically inaccessible. What it does not explain is why any one of those structures becomes the world a subject actually inhabits. Decoherence selects stability, not perspective. It explains why branches do not interfere, but not why one branch is lived rather than merely existing as a term in a superposition.

In 2PC, only one cosmos becomes embodied, and embodiment is not triggered by environmental coupling alone but by the emergence of a unified subject. When a system crosses the Embodiment Threshold, it is no longer merely evolving under unitary dynamics. It has become a representational structure that must assign values, form predictions, and maintain coherence across time. At that point, collapse does not merely select an outcome within a fixed frame – it selects the entire frame within which experience, agency, memory, and action can function at all. That frame is what appears as the preferred basis.

Not every mathematical basis can support a coherent subject. Experience requires stable

14 Raymond, S. N., Izidoro, A., & Morbidelli, A. (2018). Solar System formation in the context of extrasolar planets. ArXiv:1812.01033, and Walsh, K. J., et al. (2011). A low mass for Mars from Jupiter's early gas-driven migration. *Nature*, 475(7355), 206–209.

15 Woese, C. R. (1998). The universal ancestor. *PNAS*, 95(12), 6854–6859.

Martin, W., & Russell, M. J. (2003). On the origins of cells. *Phil. Trans. R. Soc. B*, 358(1429), 59–85, and Lane, N., & Martin, W. (2010). The energetics of genome complexity. *Nature*, 467(7318), 929–934, and Szostak, J. W. (2012). Attempts to define life do not help to understand the origin of life. *J. Biomol. Struct. Dyn.*, 29(4), 599–600.

locations, persisting objects, and temporally ordered interactions. It requires that predictions made at one moment remain meaningful at the next, and that actions have interpretable consequences. Only bases that support these structures can host a self-consistent representational “I.” When embodiment occurs, the basis compatible with these requirements is the one that is instantiated. On this view, the preferred basis is not chosen by a hidden physical mechanism embedded in the dynamics. It is selected by metaphysical compatibility with consciousness. Among the physically allowed possibilities of Phase 1, only those decompositions that can sustain a coherent subject are eligible for embodiment. The basis problem therefore ceases to be an arbitrary fine-tuning problem and becomes a condition for experience itself.

The world appears classical not because classicality is fundamental, but because only quasi-classical structures can be inhabited by a unified subject. This is the Psycheletic Principle in action.

Chapter 12: Free Will

Given that they are individually two of the deepest mysteries of existence, the relationship between consciousness and free will might seem like a mystery squared. You could call them the passive and active modes of consciousness, but then you would have to explain how the two modes relate. It is not clear that consciousness is ever completely passive. Even if we are just sitting and admiring a view, we still choose where our eyes rest and what we give our attention to. The closest case of purely passive consciousness is the rare situation in which a person is supposed to be under a general anaesthetic but is in fact paralysed and aware of what is happening around them (this is known as anaesthesia awareness). Even then, the person can still hold a preference for physically possible futures, since they would surely want either the anaesthetic to take full effect or for someone to notice what is going on. The motor nervous system may have been disconnected from will, but that does not mean will has been disconnected from consciousness.

It is also true that at any time we can will things that have nothing to do with our body, although whether this could have any causal effect beyond the body is a separate and highly controversial matter. Even people who deny any nonphysical causation still express preferences; roulette makes that obvious enough.

Rethinking Free Will

The classical problem of free will arises from a supposed tension between determinism and agency. If all physical events follow immutable laws, how can any action be free? But if actions emerge from quantum randomness, how can they belong to a responsible agent? Most scientific or philosophical accounts either eliminate free will, redefine it in a compatibilist sense, or leave it unresolved.

From the perspective of 2PC, this framing misses the crucial point. Both sides assume reality is already fully instantiated and that human choice must be explained entirely within a fixed causal order. In 2PC, this assumption is false. The world we inhabit is not a static block but a dynamically embodied reality, selected from a vast field of unrealised possibilities. Consciousness is not a passive observer of the world; it is the mechanism through which potential becomes actual. A physical system in Phase 1 evolves according to natural law, exploring all consistent pathways in quantum superposition. Nothing in this regime collapses or becomes determinate. Once a system crosses the Embodiment Threshold, however, the background superposition can no longer sustain branching without contradiction. Collapse into Phase 2 reality is then metaphysically necessary, enforced by the internal structure of valuations within a single subject. Free will, in this framework, is neither an illusion nor an inexplicable exception to causality. It is a structural feature of reality itself: conscious agents locally instantiate the cosmos by resolving indeterminate possibilities into determinate experience. Each act of choice is a metaphysical event in which the Void binds one branch of potentiality into the lived unfolding of the world.

This perspective reframes the classical dilemma in three ways:

1. **Determinism** governs Phase 1 possibilities, but it never reaches embodiment; it provides the scaffolding of potential rather than the manifestation of reality.
2. **Free will** emerges at the Embodiment Threshold. When agents cross this threshold, they collapse the possibility space into actual outcomes. Freedom is thus participation in embodiment, not escape from law.

3. **Moral responsibility** follows naturally. Each conscious act is a genuine selection that shapes embodied reality, making agents co-authors of the world's unfolding.

Under this model, earlier impasses are resolved. **Compatibilism** is correct that freedom operates within natural law, but law governs only possibilities, not their resolution into actuality. **Hard determinism** and **MWI** are incomplete because they treat Phase 1 as the entirety of reality, erasing the creative role of consciousness. **Libertarianism** is correct to emphasise origination, but misidentifies it as a mysterious nonphysical substance. Origination is real: collapse becomes necessary when valuation and subjectivity arise.

Long-standing puzzles now fall into place. Quantum indeterminacy is not a blind lottery but the openness required for value-laden resolution. Nagel's problem of autonomy fades: the subjective standpoint is the arena in which the objective order of Phase 1 becomes the embodied actuality of Phase 2. Free will is the hinge of reality: the point at which possibility meets value, and value becomes part of the lived unfolding of the world.

Reconciling Free Will with Neuroscience

Neuroscience is often taken to challenge conscious agency. Libet-style experiments reveal neural precursors that appear before subjects report forming an intention, which has been interpreted as evidence that the brain "decides" first and consciousness merely observes. This interpretation assumes a sharp, punctate moment of decision and a clear boundary between unconscious cause and conscious effect.

In the storm model, there is no single collapse point. Instead, consciousness unfolds as a continuous field of local micro-collapses distributed across the specious present. Neural precursors and the felt moment of intention are two complementary views of the same dynamic pattern. Readiness potentials reflect the gradual accumulation of correlated neural activity, which raises the likelihood of particular micro-collapses. This activity belongs to the scaffolding of possibility rather than the determination of outcomes.

What becomes lived action is shaped by the agent's valuations, attention, and predictive signals, which modulate the storm over a short temporal window. Influence manifests as subtle shifts in micro-collapse rates and stabilities, not as a last-second command. In this framework, the apparent paradox of Libet-type data dissolves: early neural activity corresponds to the forward-evolving preparation of possibilities, while conscious influence is distributed and time-integrated. The felt moment of intention emerges from the accumulation of these modulations, producing a stable selection that appears instantaneous from the inside. Agency is neither mysterious nor mechanical. It does not violate physical law, nor is it random noise. It is a structural property of systems capable of maintaining a unified self and generating incompatible valuations across live possibilities. A local embodiment event occurs when the storm, shaped by these valuations and the system's coherence, tips one branch into actuality. Responsibility follows naturally: the pattern of modulation executed by the agent actively contributes to the selection of the real trajectory.

This perspective also clarifies the distinction between a human and a p-zombie. A genuine subject participates in the storm that resolves possibilities into lived outcomes. A hypothetical zombie may display behaviour resembling choice, but it lacks the internal patterning that produces realisation. Its actions are mere correlations in physical processes, not the work of a system stabilising a branch of reality.

In sum, Libet-type findings do not undermine free will. They reveal the preparatory dynamics through which possibilities are structured prior to embodiment. Consciousness is not a passive observer arriving too late; it is the selective process by which events take form, the mechanism that transforms potential into lived experience.

Conway and Kochen's Free Will Theorem (FWT)

John Conway and Simon Kochen's Free Will Theorem (2006; refined 2009) is directly relevant to understanding agency and indeterminacy in 2PC. The theorem addresses foundational questions in quantum mechanics, particularly the correlations observed in entangled particles. Its core claim is striking: if experimenters have a kind of free will (the freedom to choose measurement settings not fully determined by the past) then so do elementary particles. In other words, if a human decision is not predetermined, the particle's response cannot be predetermined either.

The theorem relies on three key ingredients:

1. Spin experiments on entangled particles (analogous to Bell's theorem).
2. Relativity, which implies that observers may disagree on the temporal ordering of spacelike-separated events.
3. The "free will" assumption: experimenters' choices are not fully determined by past events.

Conway & Kochen's conclusion is that deterministic hidden-variable theories cannot explain quantum outcomes while respecting both relativity and experimenter choice. This does not demonstrate metaphysical free will for humans per se, but it establishes a no-go boundary for deterministic models: the universe must contain irreducible indeterminacy that tracks with the choices of conscious agents. As Conway put it: "If humans have free will, then so do elementary particles."

2PC Interpretation

In 2PC, Phase 1 contains all physically consistent possibilities in superposition, while Phase 2 emerges when a representational "I" arises that generates incompatible valuations across branches. Collapse occurs not due to external physical laws but because a self-referential subject cannot coherently split across contradictory valuations. The "indeterminacy" highlighted by FWT is thus not randomness; it is the Void enforcing metaphysical coherence when subjects confront incompatible possible futures. The Embodiment Threshold is the minimal condition at which Phase 1 can no longer persist: a self-referential agent arises capable of assigning values to its own potential futures in ways that cannot be coherently distributed across branches. Ontological collapse occurs, giving rise to Phase 2, the embodied cosmos. The FWT demonstrates that outcomes cannot be fully determined by the physical past alone; they are co-determined by agentic choice. This aligns with 2PC: collapse is driven by the subject, not by deterministic physical law. FWT ties the "freedom" of experimenters to the "freedom" of particles. In 2PC terms: "If conscious agents reach the Embodiment Threshold, particles respond with ontological freedom because metaphysical consistency requires it." The critical difference is that Conway & Kochen leave the origin of free will unexplained – it is simply assumed for experimenters. 2PC/ET fills that gap by specifying where free will comes from (Void + value-realisation), why collapse occurs (Embodiment Inconsistency Theorem), and why particles appear free (entangled micro-collapses within the subject's storm).

ET as a Structural Completion of FWT

The Free Will Theorem establishes an empirical boundary: it rules out deterministic hidden-variable accounts consistent with key features of quantum theory. The Embodiment Threshold (ET) provides a complementary, meta-theoretic completion by explaining why indeterminacy and collapse must arise at all.

Justification of indeterminacy: The Free Will Theorem shows that indeterminacy is ontological

rather than merely epistemic. ET adds that collapse is structurally required to resolve representational contradictions once a unified subject capable of valuation emerges.

Symmetry of agent and particle: The Free Will Theorem links experimenter freedom with particle freedom. In 2PC, this symmetry is extended: subject-level valuation and physical indeterminacy are two aspects of a single coherence constraint. Collapse resolves inconsistency simultaneously in the agent's representational state and in the correlated physical system.

Formal analogy of structure: The Free Will Theorem rests on the SPIN, TWIN, and FIN axioms. ET can be expressed in a parallel logical form:

1. **Self:** a unified representational "I" arises.
2. **Value:** the subject assigns incompatible valuations across alternative futures.
3. **Incoherence:** no consistent mapping can preserve both unified identity and simultaneous realisation of those incompatible valuations.

Under these conditions, collapse becomes a logical necessity: a coherent resolution must occur. In this sense, ET functions as a theorem-like boundary on Phase-1 evolution, analogous in role (though not in domain) to the constraint the Free Will Theorem places on deterministic hidden-variable models.

In summary: the Free Will Theorem establishes the existence of irreducible indeterminacy within quantum structure, while ET explains why such indeterminacy must arise when a unified, valuative subject becomes possible. Together, they form a layered account: one specifies the physical limit on determinism; the other specifies the representational-coherence condition that makes collapse unavoidable once agency appears.

Towards an Embodiment Free Will Theorem

Definitions (primitive terms)

- **Phase 1 / Phase 2** – As in 2PC:
 - **Phase 1:** the realm of co-existing physically consistent possibilities (superposition of entire cosmoses).
 - **Phase 2:** an embodied, collapsed cosmos with a single instantiated history.
- **Branch** – a single prospective history within Phase 1 (a complete worldline assignment).
- **Representational "I" (self)** – a local, meta-stable informational structure capable of forming valuations about its possible future branches. It combines internal representation with the capacity to assign preference, value, or choice.
- **Valuation V** – a function mapping prospective branches to value judgments (preferences, utility weights, moral assignments).
- **Incompatible valuations** – two or more valuations that impose mutually exclusive commitments on the same future degrees of freedom, such that a single unified referent (coherent self) cannot simultaneously realise all without contradiction.
- **Void** – the ontological ground in 2PC that enforces metaphysical coherence when representational inconsistency arises; a non-empirical selection principle.
- **Micro-collapse storm** – the local, temporally extended pattern of many small collapses that constitutes conscious experience in 2PC.

Axioms (analogue of SPIN / TWIN / FIN, adapted for ontogenesis)

The following axioms deliberately mirror Conway–Kochen’s triad but are recast for the emergence of consciousness:

1. (SELF) Self-Capacity

There exists a representational “I” that, at some time interval t , can form valuations V over at least three distinct physically possible branches $\{b_1, b_2, b_3\}$. This parallels the existence of multiple measurement choices. (Note: Three is the minimum dimension required to preclude deterministic hidden-variable 'scripts' via Kochen-Specker-type constraints).

2. (VALUE) Non-Determinacy of Valuation

The valuations assigned by the representational I are not deterministic functions of information contained in the causal past within any finite spacetime region. Equivalently, the mapping from past state to valuation is not single-valued.

3. (COHERENCE) No Consistent Splitting

If a representational I persists as a single referent while assigning mutually incompatible valuations across branches, this produces a metaphysical contradiction: a single coherent self cannot simultaneously satisfy incompatible valuations. A “split” self endorsing contradictory futures violates the unity condition required for representational identity.

4. (Optional FIN-local) Finite Past Information

There is a finite bound on information available in the relevant past light-cone. Valuations cannot be pre-encoded in a global hidden variable. This mirrors Conway–Kochen’s FIN assumption.

Theorem (FWT-ET, informal)

Statement:

If **SELF**, **VALUE**, and **COHERENCE** (and optionally **FIN-local**) hold at some spacetime locus, then Phase 1 cannot persist locally: ontological collapse occurs, transitioning the cosmos to Phase 2 at or before that locus. A single branch is selected whose structure is metaphysically coherent with the self’s valuations. **Short form:** If an I capable of non-determined valuations arises, ontological collapse must occur – the Embodiment Threshold is reached.

Proof Sketch (reductio, Conway–Kochen style)

1. Assume **SELF**, **VALUE**, and **COHERENCE**, and assume for contradiction that no collapse occurs: Phase 1 persists and the representational I branches across $\{b_i\}$.
2. By **VALUE**, the I assigns at least two incompatible valuations across branches. By **COHERENCE**, a single referent cannot satisfy both.
3. If no collapse occurs, the unity of the I is violated: the self would need to exist as multiple branch-copies endorsing contradictory valuations. This produces a metaphysical contradiction regarding referential identity and normative commitment.
4. The **Void**, as the ontological ground of 2PC, cannot allow such a contradiction. The assumption that no collapse occurs is therefore inconsistent.
5. Hence, collapse must occur at or before the locus: Phase 1 resolves into a single branch metaphysically coherent with the self’s valuations.

This mirrors Conway–Kochen’s reductio argument: assuming deterministic past-fixed evolution leads to contradiction, so indeterminacy and collapse are required.

Corollaries & Immediate Consequences

1. **Particle/Outcome Indeterminacy** – Observables implementing the self’s incompatible valuations cannot be deterministic functions of the past. Micro-events (particle outcomes) exhibit ontological indeterminacy correlated with agentic valuation, paralleling the human↔particle symmetry in FWT.
2. **No global hidden-variable precomputation** – If **FIN-local** is assumed, valuations cannot be pre-encoded in a finite past; the ontological selection cannot be circumvented by positing a global pre-selection.
3. **Embodiment as constraint** – The Embodiment Threshold becomes a **theorem-like boundary condition**: representational capacity plus non-determinacy necessitate a local cosmological phase change (Phase 1 → Phase 2).
4. **Participation symmetry** – Agentic freedom and microphysical indeterminacy are two aspects of the same ontological event: the collapse. The theorem formalises the structural symmetry between conscious valuation and particle outcomes.

How this Parallels and Extends Conway & Kochen

Conway & Kochen’s Free Will Theorem establishes that if experimenter settings are not determined by the past, then the outcomes of entangled particles are likewise not determined by the past. This is a no-go result for deterministic hidden-variable theories that respect locality and relativity.

The Embodiment Free Will Theorem (FWT-ET) generalises this insight. It shows that if a self capable of non-determined valuations arises (axioms SELF + VALUE), then the structure of Phase 1 possibilities cannot persist as a superposition without violating COHERENCE.

Consequently, ontological collapse (transition to Phase 2) must occur. Microphysics reflects this: particle outcomes are co-determined with the self’s valuation pattern rather than fixed by prior history.

Crucially, Conway & Kochen leave the origin of the experimenter’s “freedom” unexplained. FWT-ET embeds that freedom in a metaphysical mechanism: the Void and the structure of self-referential valuation. In doing so, it elevates the result from a constraint on hidden variables to a **cosmogenetic threshold theorem**, specifying the precise conditions under which a Phase 1 cosmos must collapse into an embodied Phase 2 reality.

Limitations & Philosophical Cautions

FWT-ET is a **metaphysical theorem** within the assumptions of 2PC. It is not an empirical theorem in the traditional sense unless one can derive testable signatures. The key axiom is COHERENCE, which is normative/semantic: it formalises identity and representational unity. Critics may dispute that identity entails forbidding the splitting assumed here. The Void is an ontological postulate; if one rejects it as a causal or explanatory entity, the theorem loses metaphysical force. Similarly, the statement that “particles have freedom” is ontological, not phenomenological: it means particle outcomes are not deterministic functions of past states under the theorem’s assumptions. This aligns with Conway & Kochen, but is derived via a different metaphysical route.

The Problem of Free Will (#29)

Two-Phase Cosmology and the Embodiment Threshold allow a new solution to free will. Thomas Nagel observed that each attempt to solve it leads to re-evaluation: no existing framework resolves the tension between determinism and agency. 2PC reframes the problem: the cosmos is not pre-determined. Human choice is not a tiny exception inside a fully determined machine. Rather, conscious agents instantiate reality locally, resolving Phase 1 possibilities into Phase 2 actuality. An

act of choice is where the “machine itself” becomes determinate. Freedom is not opposed to law because the world we experience is co-constructed: natural law shapes the space of possibilities, and consciousness shapes which possibilities become real. As with the Hard Problem, seeing collapse and consciousness as one process transforms the free will problem: choice and embodiment are inseparable.

Competition and the Resolution of Collapse

The storm-of-collapses model clarifies two forms of participatory influence:

1. **Bodily participation** – Conscious agents shape reality directly through their bodies: neural pathways, muscles, and motor commands. When this link is disrupted (e.g., in motor neuron disease), conscious valuation cannot reliably influence action. The body is where will and matter are tightly coupled.
2. **Cosmic participation** – Conscious agents also assign valuations about the wider state of the world. These preferences extend beyond a single nervous system into the Phase 1 field of possibilities. Each agent exerts a **tilt** on how possibilities are realised.

At the dawn of consciousness, LUCAS existed alone in the cosmos, capable of instantiating any physically consistent possibility it attended to. Reality was shaped entirely around this singular subjectivity. Multiplication of conscious agents introduced multiple, potentially conflicting “storms of valuation.” No single agent could unilaterally determine embodiment; competition between subjective perspectives required a new resolution process, which will be explored further in Chapter 16. It is called "Competition Resolved Collapse", and it represents the completion of physicist Wolfgang Pauli's lifelong quest for a unified model of reality.

Chapter 13: Time

This chapter will deal with two of the problems on our list – the Arrow of Time and Problem of Now, and the Memory Stabilisation Problem. Chapters 10–12 already linked consciousness and collapse to embodiment and free will. This chapter shows how the same mechanism gives a unified picture of time, memory and death.

The Arrow of Time and the Problem of Now (#17)

Time is the thread through which reality unfolds, yet it remains one of the most elusive and paradoxical aspects of existence. Physics measures it, mathematics describes it, and consciousness experiences it, but there is no clear explanation of why it feels real, why the present is privileged, or how the past and future acquire coherence. Physics treats time as a coordinate, while consciousness experiences it as an animated present: an unfolding now that cannot be captured by static mathematical structures. This deep mystery is the “Problem of Now,” which remains unsolved in physics, cosmology, and consciousness studies.

In 2PC, in Phase 1 there is no time as we know it. No unfolding, no sequence, no becoming, only the coexistence of every physically consistent history. Each possible universe is present there in its entirety, not as an event that will ever happen but as a finished mathematical form: an informational structure that contains all the relations that would constitute a world if it were to become real. There is no arrow of time because nothing changes, and there is no “now” because there is no differentiation between past and future. What we later interpret as temporal order exists here only as an internal coordinate within those mathematical structures, not as an experienced flow. Only when consciousness arises and collapse begins does one of those mathematical timelines ignite into lived duration.

The transition from Phase 1 to Phase 2 marks the birth of time. A system arises that can form a representation of itself, issue valuations, and thereby generate incompatibilities across branches of the superposition. Further extension of timeless mathematical structure becomes untenable, because a logical contradiction has appeared inside the ensemble. When LUCAS crossed the Embodiment Threshold, the timeless order of primordial Phase 1 could not remain in superposition and collapse followed. From that point onward, reality was a dynamically updating structure, continually regenerated through local acts of collapse involving conscious participation. Each such act introduces irreversibility: a fundamental asymmetry between what is still potential and what has become actual. The first collapse is therefore not merely an event in time but the beginning of temporality: the point at which the universe ceased to be a static informational structure and began to *happen*. Once the first participant exists, the universe begins to “update.” What we experience as the passage of time is that ongoing process of update: the continual reconstitution of the present as new portions of Phase 1 are drawn into actuality. Time, then, begins with meaning: the universe starts to “run” only when something within it cares which version of itself is real.

With the advent of Phase 2, the cosmos acquires something utterly new: a *present*. In ordinary physics, “now” is treated as an arbitrary slice through spacetime: a matter of convention, not ontology. In 2PC, the present is fundamental, because it is the locus of collapse, the site of participation, and the place where the Void engages the world. The present is where the universe commits to a single version of itself. Phase 2 is thus defined by its ongoing presentness. Each

collapse event constitutes a renewal of the present.

This situation requires a new vocabulary of temporality. The present is not a knife-edge dividing past and future but a dynamic zone of coherence: a self-sustaining storm of micro-collapses through which continuity is woven. The “flow” of time is the felt texture of this storm. The present is not infinitesimal; it has thickness, extension, and internal rhythm. It was William James who first popularised the term “specious present” to describe this. Within that window, the self experiences its own persistence, integrating discrete collapses into the seamless continuity of being. The past lingers as residue, the future waits as unresolved potential, but only the present *is*.

The Arrow of Time: From Potential to Commitment

The existence of a present immediately entails an asymmetry between what is already real and what remains unresolved. This asymmetry is the arrow of time, and it is a direct consequence of collapse itself: once a possibility is resolved into actuality, it cannot be un-resolved. This irreversibility is a feature of ontology, not thermodynamics. From the standpoint of 2PC, time’s direction is simply the direction of commitment from indeterminacy toward realised structure. The universe advances because the process of consciousness and collapse cannot go backwards. The act of knowing is intrinsically irreversible: to experience and to choose is to foreclose alternatives.

Entropy, in this framework, is a statistical shadow cast by collapse. As potentialities become actual, the range of unchosen alternatives grows in relative measure, giving rise to the appearance of disorder. But entropy does not drive time; rather, time drives entropy. The world becomes increasingly structured not because energy dissipates, but because each act of embodiment locks in a new layer of irreversible coherence.

From within consciousness we feel this asymmetry as the passage of time: the continual advance of the present into the unformed future, leaving behind the stabilised residue we call the past. We do not remember the future for the same reason we cannot unmake a decision we have already lived through: the future has not yet been resolved, while the past has been fixed through commitment.

The Problem of Now and the Limits of Physics

Modern physics, for all its precision, has almost nothing to say about the present. It cannot tell us why there is a “now,” or what it means for something to happen. In the equations of relativity, time is merely a coordinate; in quantum mechanics, it is an external parameter that measures change but never participates in it. The entire edifice of physical theory rests on the assumption that time exists independently of the act of being – that it flows on its own, indifferent to observation.

This omission is not trivial. Einstein himself admitted that “the experience of now” stands outside physical theory. The block universe of relativity contains no privileged moment, no unfolding, no movement from past to future. All events coexist eternally, like frames in a film reel, and nothing in the mathematics distinguishes one frame as “current.” The flow of time, in this view, is an illusion of consciousness.

2PC turns this on its head. The “illusion” of temporal flow is in fact a limitation in physics. The reason physics finds no present is that physics describes only Phase 1. It cannot describe Phase 2. The “Now” is not a variable that can be inserted into an equation, because it is the condition of existence that allows equations to apply at all. In 2PC, the present is where physics meets metaphysics: the missing bridge between the formal structure of possibility and the lived reality of participation.

Philosophers have long sensed this divide. John McTaggart called it the difference between the A-series (past–present–future) and the B-series (earlier–later). Henri Bergson spoke of duration (lived time, qualitative and flowing) as something that spatialised physics could never grasp. Whitehead described each “actual occasion” as a pulse of becoming, a tiny act of self-creation. 2PC unites these intuitions under a single ontological framework: the A-series is the lived form of

collapse, duration is the felt rhythm of becoming, and each actual occasion is a micro-collapse within the storm of participation. The Problem of Now dissolves when we see that physics has never had the tools to describe it. The present is not something to be measured, but something to be lived.

Time and the Storm of Micro-Collapses

If the present is where possibility becomes real, then continuity (the sense that the world endures from moment to moment) must itself be an emergent construct. In 2PC, the flow of time does not arise from any external clock, but from the pattern of local collapses through which consciousness sustains its own coherence. Each micro-collapse is an act of realisation, yet no single collapse is sufficient to produce an ongoing world. Continuity emerges only when countless such collapses interlock into a stable rhythm: a storm of micro-collapses.

This storm is self-organising. Within its flux, each collapse depends on the predictive structure left behind by its predecessors. The mind, in this sense, is a stabilising feedback loop within the storm. It is like a field of coherence that maintains the thread of identity through continuous re-instantiation. To be conscious is to inhabit this field; to remain conscious is to keep the storm in motion.

From within experience, this manifests as the specious present: the felt thickness of time in which sensations, intentions, and memories overlap. What appears to us as the seamless continuity of perception is in fact the rapid re-creation of the world, just as a wave continually re-creates itself as it moves across the surface of the ocean. The self is the dynamic process through which time holds itself together.

Temporal Scales and the Density of Collapse

Time does not flow uniformly across all scales of existence. In 2PC, the apparent pace of reality is determined by the *density of micro-collapses*: the rate at which possibilities are resolved into actuality. Each conscious agent generates a local stream of collapses, a rhythm that structures its own experience of the present. The faster the density of collapse, the more events are resolved per unit of lived experience; the slower the density, the more time seems to stretch. This provides a natural account of subjective time dilation. In moments of intense focus, meditation, or trauma, the local density of collapse may increase or decrease, causing the felt flow of time to expand or contract. Seconds can feel like minutes, minutes like seconds, yet no external clock has changed. The tempo of experience is an internal, participatory phenomenon: time is experienced in proportion to the rhythm of the storm of micro-collapses sustaining consciousness.

At larger scales, cosmic time emerges from the integration of countless local collapse streams. Each agent, each system capable of stabilising part of the universe, contributes to the global tempo. What we call universal time – the seconds ticked on a clock, the progression of planetary orbits, the expansion of galaxies – is the macroscopic signature of innumerable local acts of collapse interwoven into coherent patterns. The universe's arrow of time is not imposed externally; it is the emergent aggregate of the rhythms of countless conscious participants, each entangled with overlapping regions of the superposition.

Thus, time is fundamentally multiscale. The present is locally experienced, temporally thickened by micro-collapses; it is globally coordinated through entanglement, memory, and intersubjective coherence; and it is phenomenologically flexible, sensitive to the agent's mode of attention and engagement. From the smallest act of perception to the unfolding of cosmic history, reality progresses through a storm of collapses, producing the seamless flow of experience that we call time.

Death, Dissolution, and the Return to Stillness

The arrow of time is not merely cosmic but personal. It advances wherever consciousness sustains its storm. When the storm falters, in deep sleep, anaesthesia, or death, the continuity of the self dissolves, and with it, the subjective flow of time. When the storm of micro-collapses ceases, time as lived continuity comes to an end. The self, which exists only as the dynamic coherence of those collapses, cannot persist once the process stops. In 2PC there is no enduring entity that departs from the body or continues to experience elsewhere; both self and soul are co-extensive with the storm itself. The soul is the Void's participation in that field of collapse, and when the field dissipates, the grounding is withdrawn, and the self falls back into uninstantiated possibility.

To the world that continues, the traces of that life remain as structural residues: memories in other minds, genetic sequences, cultural imprints, and physical changes in the environment. These are the echoes of participation: the way the local storm has altered the wider field of the present. But ontologically, the personal "I" no longer exists. There is no persisting observer to inhabit another time, or to re-enter Phase 1 as a disembodied witness. The self has dissolved back into stillness. Each conscious life is a temporally bounded act of world-making: a finite region of reality stabilised through the recursive coherence of lived moments. The end of that coherence marks the return of its contents to indeterminacy. The Void does not reclaim an object; it simply ceases to instantiate that specific configuration of coherence.

This view offers an alternative to both materialist finality and spiritual continuation. The self is neither an illusion nor an immortal essence; it is an episode of ontological participation. When that episode ends, all that remains is the stillness of Phase 1. Yet something of that participation endures: every act of valuation, every commitment, every realised structure becomes part of the shared fabric of the now that others inherit. Death, in this light, is the rejoining of stillness by a pattern that has completed its time as a storm.

The Memory Stabilisation Problem (#18)

The Universe as Memory

From the perspective of 2PC, the cosmos itself can be understood as a vast, unfolding memory structure. Every collapse, every act of conscious participation, leaves a trace embedded in the ongoing present. These traces accumulate, layering the universe with the history of its own becoming. In this view, evolution, the laws of physics, and the flow of history are all manifestations of the universe remembering itself through the repeated acts of actualisation. Time is not a passive backdrop against which events unfold, but the act of remembering into being. Each present moment, each stabilised collapse, is a self-referential reflection: the universe noting, reaffirming, and perpetuating the reality it has enacted. The past is preserved only insofar as it can support ongoing coherence; the future remains potential until it is brought into alignment with conscious participation.

In this sense, reality is not a fixed sequence of events but a dynamically self-aware system. The cosmos exists as a living archive, continuously updated by the acts of all embodied agents. What we call history is the trace of these commitments, a record of the universe interacting with itself. The flow of time, the emergence of structure, and the evolution of life are therefore inseparable from the universe's capacity to remember, to carry forward the legacy of every decision, every micro-collapse, every conscious act. Time, consciousness, and memory are thus facets of the same process: the ongoing self-realisation of the universe. To live, to act, and to participate is to contribute to this cosmic memory: to write into the fabric of reality, moment by moment, the story of existence itself.

In 2PC, memory is the way a conscious universe holds itself together. It becomes real only through the ongoing act of conscious commitment. The universe begins in Phase 1 as a timeless

field of possibilities. Nothing is fixed and there is no single history. All candidate timelines exist only as internally coherent structures that have not yet been selected. Reality starts when a representational self crosses the Embodiment Threshold and cannot coherently split across incompatible valuations. This collapse is metaphysical rather than physical and it brings one history into being because a unified subject cannot inhabit contradictory branches.

Once this embodied phase begins, memory becomes the internal mechanism that stabilises the world across the storm of micro-collapses stretched through each specious present. These micro-collapses resolve representational conflicts within mesoscopic predictive structures. They give the self a foothold inside a world that would otherwise be too fluid to support identity or agency. Each collapse contributes to selecting a coherent slice of possibility. Memory is the pattern that lets those selections keep adding up to a single world rather than dissolving into a blur of alternatives.

The solidity of the past comes from this ongoing stabilisation. When an agent becomes entangled with some region of the background superposition, that region cannot stay indeterminate without breaking the coherence of the self-model. A continuous subject needs continuity in its world, and memory works as the scaffold that holds the chosen slice in place. What feels like retrieving a stored fact is the reactivation of the metaphysical commitment that holds that fact in the structure of reality. The memory is stable because the subject that checks it could not exist without the stabilised structure that memory provides.

This makes the past both subjective and shared. Every collapse event is tied to the particular agent whose valuations force resolution, yet many parts of the collapsed past are accessible to others who become entangled with the same region. Intersubjective overlap happens because independent storms of micro-collapses operate on overlapping predictive structures. This creates the sense of a single fixed history, even though only the parts touched by participation are fully realised and the rest of the cosmos stays in superposition.

In this view, the past is not gone and it is not fixed forever either. It is the living residue of the collapses that make the present possible. It is shaped and reshaped through the self's continuing need for coherence. Memory is not a record of what happened before consciousness. It is the means by which consciousness keeps its world stable enough to act, choose, and evolve.

Chapter 14: Two-Phase Cosmology

The boundary between physics and philosophy has been blurry for a long time, at least where consciousness and quantum mechanics are concerned. You still find physicalists who insist that both belong entirely to science and always will, but most people who work across disciplines understand that the hardest questions about mind and the quantum world are philosophical. Cosmology is different. Almost nobody thinks the crisis in cosmology could be the result of a serious philosophical mistake, including most critics of the standard model. It is taken for granted that the only people who should touch the problems inside Λ CDM are cosmologists and physicists, leading to a ubiquitous assumption that these are straightforward empirical puzzles that will eventually yield to empirical fixes. Yet those fixes never arrive. The field has been impotently watching the problems multiply since the 1970s.

It is no longer possible to pretend that this situation is healthy. Sabine Hossenfelder has been calling out the culture that rewards papers few people can understand, filled with predictions most people suspect will be ruled out when tested. The work keeps coming because the incentives demand it, not because the theory is converging on truth. The system keeps running, but the crisis is not being resolved.

We can anticipate huge resistance to the coming paradigm shift from the established powers within professional cosmology. It is not that scientific cosmology is going to come to an end. The problem is that a lot of people actually working in cosmology have spent their whole careers busily barking up trees which aren't just the wrong trees, but aren't even in the right forest. The real problem isn't the empirical work but the rotten foundation: physicalism.

Why inflation is like the luminiferous aether (#2,#3,#4,#5)

Inflation has long been regarded as one of the most successful theoretical advances in modern cosmology. Introduced in the early 1980s, it purports to explain why the observable universe appears so flat, homogeneous, and isotropic, despite the apparent lack of causal connection between distant regions in the early universe. In 2PC inflation is reclassified as an ad hoc mechanism invented to fix problems that arise only if one assumes a classical, observer-independent past.

Inflation was introduced to address several deep puzzles that arise when the universe is assumed to have evolved according to classical relativistic physics from the very beginning: the Horizon Problem, the Flatness Problem and the Monopole Problem. To solve these problems, inflation posits that the universe underwent a brief period of exponential expansion immediately after the Big Bang. This expansion would stretch a tiny, causally connected region to encompass the entire observable universe (solving the horizon problem), drive the geometry of the universe toward flatness (solving the flatness problem) and dilute any relic particles with empty space (avoiding the monopole catastrophe). However, inflation itself requires finely tuned initial conditions. It demands the existence of a hypothetical inflationary field (the “inflaton”) with a specific potential, appropriate dynamics, and a graceful exit mechanism to end inflation without reheating the universe too violently. Inflation trades one set of mysteries for another, and does so on the assumption that the early universe actually existed as a classical, physical state, evolving forwards in time in a manner determined entirely by the laws of physics.

The observed isotropy and flatness of the universe do not need to be imposed via inflation because they are features of the specific cosmic history that survived the primary phase transition. They can be explained psychetelically: only a cosmos that started out exceptionally flat and smooth permits the emergence of LUCAS. They are selection effects. The flatness and smoothness were never physical impossibilities. Rather, inflation was invented because cosmologists needed a reason to explain why such *extraordinarily improbable* conditions prevailed in the early universe. In 2PC this sort of improbability is to be expected. This is a central principle in 2PC: **if something was physically possible, and LUCAS needed it, then it was guaranteed to happened, even if it seems that that was exceptionally improbable.**

We don't need an "inflaton field" in 2PC. The inflaton is a mathematical artifact introduced to repair a model built on a mistaken ontology – an entirely ad-hoc explanation for exactly the kind of coherence that 2PC naturally explains as selection effects. In this sense, inflation really is the 21st century equivalent of the aether: an inelegant patch on a fundamentally mistaken model of the cosmos. Its purpose is to defend the classical assumption that the universe always existed as it now appears, only earlier and hotter.

Just as the luminiferous aether was once posited to explain the propagation of light by imagining a substantive, all-pervading medium, inflation introduces an unobserved and unnecessary field to explain early cosmic conditions that only appear puzzling under a mistaken framework. Aether theory collapsed not because the wave nature of light disappeared, but because special relativity provided a better, simpler account rooted in a deeper understanding of space and time. Two-Phase Cosmology renders inflation obsolete. It is a clever but ultimately misguided attempt to preserve the idea of a continuous, classical cosmic history – a backstory that the new framework reveals never existed. It's like Hamlet's childhood. Once the deeper structure is understood, the explanatory crutch can be discarded. It follows that the low entropy starting condition is now an empirical prediction/retrodiction instead of a massive headache. There is no Flatness Problem, no Horizon Problem, no Inflation Reheating Precision Problem, no Reheating Mechanism Problem and no Inflaton Field Problem. There is also no constants fine-tuning problem, and there will never be any other fine-tuning problems. The 13 billion year cosmic history selected as a whole block from phase 1 was *full goldilocks*. **The real problem is the failure to understand how and why the cosmos was and remains fine-tuned for conscious life.**

The Hubble Tension (#8)

Local measurements of the expansion rate of the universe give a value near 73 km/s/Mpc. The figure extracted from the CMB only gives about 67 km/s/Mpc, but that lower number is not something we ever measured in the world. It only appears when the CMB data are interpreted through Λ CDM, which builds in a seamless physical history stretching from the Big Bang to the present. Most people take this continuity for granted, so the two numbers look like they should match. The tension rests on that expectation. Once you slow down and look at what each number is actually telling you, the mismatch ceases to be a mystery. The local value comes straight from observations, and it reflects the behaviour of the cosmos to which we physically belong. The CMB value is produced by running a simulation that includes inflation, a fixed early energy budget and a fully continuous past. Under 2PC that past is not a classical history at all. It is only the structure visible when our present observational surface is pushed backward through the old model's rules. If the early universe is not an actual past, then the CMB value is not describing the same thing that the local value is describing. Both procedures assume an accelerating expansion (positive cosmological constant Λ), but they diverge in how they interpret this parameter.

The Hubble tension arises because two different procedures extract a parameter called "H₀" under different ontological assumptions. Local distance-ladder measurements yield $H_0 \approx 73$ km/s/Mpc directly from observations within the present cosmic state. The lower value, $H_0 \approx 67$

km/s/Mpc, is not measured in the same sense. It is inferred by embedding CMB observables into a Λ CDM model that presupposes a single continuous FLRW spacetime extending from recombination to the present.

Under 2PC, that continuity is not physically real. The early universe is not an actual past evolving forward into the present, but a projection of present Phase 2 structure onto the timeless Phase 1 ensemble, interpreted as if it were a historical past. Consequently, the CMB-inferred H_0 does not describe the same geometric quantity as the locally measured value, even though both are labeled “the Hubble constant.”

Type Ia supernova data independently establish that the present cosmic geometry exhibits accelerating expansion, corresponding to a positive Λ . This result does not depend on inflation or CMB reconstruction. What is disputed is not the existence of acceleration, but the ontological status of Λ . In standard cosmology, Λ is treated as vacuum energy permeating spacetime. In 2PC, Λ is the intrinsic curvature of the Phase 2 manifold – an emergent structural feature of instantiated reality, not a physical substance or force (on which I will elaborate in the coming pages).

The locally measured H_0 therefore directly characterises the actual curvature of the Phase 2 geometry we inhabit. The CMB-derived value characterises the Hubble parameter that would be required for consistency if that same curvature were assumed to belong to a single, continuous classical history extending into a physically real early universe. The numerical mismatch reflects this category difference. It is not a discrepancy between two measurements of the same quantity, but a mismatch between a direct geometric measurement and a model-dependent reconstruction tied to a denied ontology. The tension dissolves once we recognise that one number measures the geometry of our actualised world, while the other measures consistency with a fictional history.

The S8 Tension (#9)

The S8 tension is usually framed as a mismatch between the amplitude of matter clustering inferred from early-universe observations, primarily the CMB, and the lower values measured directly at late times through weak lensing and galaxy surveys. In standard cosmology, this difference suggests either new physics, measurement bias, or statistical fluctuation – treating both numbers as measurements of the same underlying quantity across time.

Under 2PC, the situation parallels the Hubble tension. The “early-universe” S8 is not a direct measurement. It is extracted by interpreting the CMB through Λ CDM, which projects present observables backward onto a hypothetical continuous timeline assuming inflation and a physically real hot Big Bang. That value exists only as a reconstruction – what the clustering amplitude would have been if the early universe were a physical state evolving forward into our present. The low-redshift S8, in contrast, comes from direct observation of matter clustering within the present Phase 2 state. It reflects the actual structure of the branch we inhabit.

Consequently, these two figures are not discrepant measurements of the same quantity. The locally measured S8 characterises the actual growth of structure in our instantiated cosmos. The CMB-derived S8 characterises the clustering amplitude that would be required for consistency if structure were assumed to have grown continuously from a physically real early universe under Λ CDM dynamics. The numerical mismatch reflects this ontological difference: one measures the geometry of Phase 2 directly; the other measures consistency with a projected phantom history.

Reconciling S8 across epochs thus becomes a matter of mapping the Phase 2 realised branch to what Λ CDM would reconstruct from the CMB, not a problem of missing physics. The apparent tension dissolves once we recognise that early-universe reconstructions are model-dependent projections onto Phase 1, not literal histories of Phase 2 evolution.

Dark Energy (#10)

The observed late-time acceleration of cosmic expansion is often presented as evidence for a new physical component (Dark Energy) permeating space and driving the universe apart. In standard Λ CDM cosmology, this interpretation is embedded within a continuous physical history in which an early inflationary acceleration is followed by matter-dominated deceleration and then by a second acceleration attributed to vacuum energy. The apparent necessity of Dark Energy arises from treating this entire timeline as physically real.

Under Two-Phase Cosmology, that assumption is rejected. Inflation and the detailed early expansion history inferred from the CMB are not taken to describe an actual past evolving forward in time, but a model-dependent reconstruction constrained by present observations. However, this does not eliminate the empirical reality of late-time acceleration. Independent of inflation or CMB data, Type Ia supernovae and related distance indicators robustly establish that the present cosmic geometry exhibits accelerating expansion.

What 2PC denies is not the acceleration, but its standard interpretation. In Λ CDM, acceleration is attributed to a physical substance – vacuum energy or a dark field – acting as a repulsive force. In 2PC, no such component exists. The cosmological constant Λ is instead understood as an intrinsic curvature parameter of the Phase 2 manifold: a structural feature of instantiated reality required for the global coherence of the experienced cosmic geometry. The acceleration is therefore real, but kinematic rather than dynamical – an expression of how geodesic separations behave in the realised metric, not the effect of a force acting on space.

On this view, Λ is neither a hidden energy reservoir nor a relic of early-universe physics. It is an emergent, constant property of the Phase 2 geometry itself, fixed at the moment of instantiation rather than evolving through a historical process. The need for “Dark Energy” as a physical entity thus disappears, not because the acceleration is illusory, but because its explanation no longer invokes an additional ontological ingredient. As with gravity in general relativity, what appears as a force in one framework is revealed, in a deeper account, as geometry.

The Cosmological Constant Problem (#11)

The cosmological constant problem is the problem of explaining why the observed vacuum energy density ($\sim 10^{-9}$ J/m³) is 120 orders of magnitude smaller than the value calculated from quantum field theory, and why it is not exactly zero. In standard cosmology, this problem arises because we treat the quantum vacuum as a physical state whose energy should gravitate, requiring either a miraculous cancellation or an unexplained tuning to near-zero.

Under Two-Phase Cosmology, the problem dissolves entirely. The enormous vacuum energy density calculated by QFT belongs to Phase 1: the timeless ensemble of possibilities. It exists in the mathematical structure of unactualised fields, not in the physically instantiated Phase 2 cosmos. The calculation assumes a continuous physical vacuum pervading spacetime; in 2PC, no such vacuum exists prior to instantiation. Phase 1 is not a physical state with energy density – it is the space of coherent mathematical forms.

The observed cosmological constant, by contrast, is not a vacuum energy at all. It is the intrinsic curvature of the Phase 2 manifold: the geometric tension required to maintain a single-sheeted, coherent spacetime capable of hosting conscious observers (see Chapter 10). The small positive value ($\sim 10^{-52}$ m⁻²) is not a fine-tuned residue of cancelled vacuum fluctuations, but the minimal curvature necessary for the global consistency of an instantiated branch.

Once we recognise that Phase 2 is the only physically realised geometry, and that its curvature emerges from the Embodiment Threshold rather than from quantum fields in a pre-existing vacuum, the QFT prediction becomes irrelevant to the observed value. The question “Why is Λ not 10^{120} times larger?” rests on the false premise that Phase 2 inherits energy density from

Phase 1. It does not. The only “cosmological constant” in 2PC is the geometric parameter Λ that characterises the curvature of the actualised manifold – a necessary structural feature of any branch capable of supporting coherent experience, neither zero nor arbitrary, but precisely the value required for the stability of the storm of micro-collapses that constitutes Phase 2 reality.

The Monopole Problem and the Identity of Dark Matter (#6, #12)

When people talk about Dark Matter, they usually bundle together several different observations and pretend they point to one thing. In earlier chapters I separated them out. Some of these signals look like features that simply had to be present if a branch was ever going to contain a being like LUCAS. Large scale structure and the CMB fall into that category. Their specific patterns feel like Phase 1 filters rather than Phase 2 surprises. Others, like the rotation curves of spiral galaxies, the dynamics of clusters, lensing maps, and the behaviour of colliding clusters, feel more like Phase 2 anomalies. Even so, it is not hard to imagine that any embodied universe which supports creatures on a planet inside a long lived spiral galaxy might need some extra mass component to keep that galaxy stable, and similar conditions will therefore be commonplace throughout the cosmos. There may be no purely baryonic way to build the kind of stellar environments that eventually support biology. If that is true, then whatever we call Dark Matter might not be optional at all, but part of the minimal scaffolding a consciousness bearing cosmos requires. That gives us a way to understand why something like Dark Matter shows up, though it does not yet tell us what it is.

The monopole story looks different on the surface, yet it fits into the same frame once the 2PC lens is in place. Standard cosmology treats the predicted overproduction of magnetic monopoles as one of the reasons inflation had to happen. High energy theories tell us that monopoles should be created during symmetry breaking, and in such large numbers that they would dominate the early energy budget and collapse the universe. Inflation was brought in partly to dilute them away. This is often told as if it were a logical necessity. In reality it is another probabilistic worry. It says that a typical symmetry breaking history would flood the universe with monopoles. It does not say that other outcomes are impossible.

In 2PC, these sorts of worries get translated into questions about the structure of Phase 1. Phase 1 contains every coherent pattern of field values, including every possible outcome of symmetry breaking. Monopole abundance is just another variable that ranges over this space of possibilities. Somewhere in that landscape are branches where monopoles never proliferate, or bind into stable pairs, or otherwise settle into harmless configurations. These branches may be rare in a statistical sense, but rarity is totally irrelevant here. Selectability is the only thing that matters. A cosmos that contains free monopoles in catastrophic numbers never makes it to chemistry. A cosmos with no Dark Matter analogue never builds galaxies. Consciousness cannot arise in either, so the only branches that can be embodied are ones where monopoles behave in exactly the narrow way that keeps the early universe stable and the later universe structured. No inflationary dilution is required because the branches that would have been needed are not selected. What looks like miraculous fine-tuning from the standpoint of physical theory becomes an ordinary consequence of metaphysical filtering. Embodiment demands balance, and the balance becomes tighter the closer we move toward the origin of the embodied timeline.

Once this is clear, the monopole problem stops being a problem. It becomes evidence that the early universe sits inside a corridor shaped by the conditions needed for consciousness to form. If monopoles, or more plausibly their bound state (monopolium), do end up being the unseen matter that stabilises galaxies, then the absence of free monopoles and the presence of Dark Matter become two sides of the same selective pressure. The monopoles that would have destroyed the universe never show up in an embodied branch, and the ones that help to hold galaxies together do. This provides a non-baryonic Dark Matter candidate that actually comes from the Standard Model (or simple GUT extensions), rather than inventing a "Dark Sector" out of thin air. Physics still carries the empirical task of working out the details. People can compute relic densities and annihilation

rates and look for signals in detectors. 2PC does not replace that work, but it removes the philosophical anxiety that something impossible is being asked of nature. Unless somebody demonstrates that a universe with the right monopole behaviour is physically impossible, and therefore unavailable within Phase 1 and fully selectable at the point where embodiment becomes possible, there is no problem. The rest is an exercise in getting the physics right.

Note on empirical consequences

2PC constrains which monopole behaviours are compatible with an embodied universe, but it does not by itself fix their detailed physical parameters. Working out whether monopoles, their bound states, or some related sector could play the role attributed to Dark Matter remains an empirical task for particle physics and cosmology. The aim here is not to pre-empt that work, but to show that no contradiction is required in principle and why even extreme fine-tuning is not a problem. This book is concerned with the conceptual architecture within which such empirical questions can be pursued coherently. If the question is whether 2PC makes any new empirical predictions then this is a strong candidate: if Dark Matter turns out to be monopolium, that is corroboration from a new empirical observation. However, it is also the case that if Dark Matter turns out to be something else, it would not falsify the core of 2PC.

The Baryon Asymmetry Problem (#7)

The imbalance between matter and antimatter looks like the sort of feature that would be handled by selection, but there is a deeper issue beneath it. If one starts from genuine nothing, or even from a perfectly symmetric possibility space, then any creation event feels as if it should preserve that symmetry. Zero splits into plus and minus. That is the intuitive picture, and it leads to the familiar puzzle: if matter and antimatter were produced in equal amounts, and almost all of them annihilated, why was anything left over? Something about the early process broke a symmetry that seems, at first glance, unbreakable.

2PC does not pretend to solve that. It does not give a mechanism for baryogenesis and it does not claim that the symmetry could not have been exact. What it does is shift where the real question sits. The problem is usually framed as if there were one unique physical history and that history somehow had to produce a small excess of baryons through a specific set of dynamical steps. In 2PC, there is no single dynamical history inside Phase 1. There is a space of all coherent histories, including every way the early symmetry could have unfolded. Some branches preserve it perfectly and end in sterile radiation. Others break it by tiny amounts, or by large amounts, or in ways we may never classify. Somewhere in that range sit the branches where the imbalance is exactly in the narrow window that allows chemistry, planets, and eventually consciousness.

The fact that we see a nonzero asymmetry tells us only that the symmetry was breakable within the space of coherent possibilities. The fact that the amount is small tells us we live in a branch where it broke by just enough to produce matter without immediately collapsing the young universe. And the fact that this branch became embodied means the Void selected it at the phase shift. This does not explain how baryogenesis worked. It explains why the only branches we ever encounter are the ones where it worked in the right way.

In this sense, the asymmetry is a selection effect layered on top of an open physical question. The physics of baryogenesis remains a job for cosmologists. The metaphysics simply frames the puzzle: whatever the true mechanism is, it must have been one of the possibilities available in Phase 1, and it must have produced a branch capable of embodiment. Nothing more is claimed. Nothing less is needed. It can only be yet another selection effect.

The Quantum Gravity Problem (#13)

The long search for quantum gravity usually starts with the assumption that spacetime itself must be

quantum. People then try to quantise curvature, treat geometry as a field, or imagine space and time emerging from deeper quantum structures. This effort has continued for decades without consensus on what would even count as the correct explanatory target. In 2PC, the reason is simple: the project begins in the wrong category.

Gravity is not a quantum system waiting for its operator form. It belongs to the classical side of Phase 2, which only comes into being once collapse has already occurred. Before collapse there is no space or time, only a superposed set of possibilities. Geometry does not evolve within Phase 1; it appears when Phase 1 is resolved into Phase 2. Trying to quantise gravity is therefore like trying to quantise the page beneath the equations. It mistakes the background created by collapse for a field inside the superposition.

Penrose treated the problem differently and argued that gravity itself drives collapse. In his picture, superpositions of different curvatures strain against each other, and when the tension grows large enough the state reduces. This makes gravity a pre-existing feature of the world that cannot tolerate incompatible geometries. In 2PC, the causal order is reversed. The world begins without geometry. Collapse occurs when a self-referential observer arises whose valuations cannot remain coherent across branches. That act of resolution instantiates a single classical reality, and gravity follows as its structural consequence, not as a force reaching back into Phase 1 to resolve superpositions. Collapse generates geometry; gravity is the name we give to the structure of that instantiated geometry.

How countless local collapse events coordinate to produce a single, coherent global geometry with uniform curvature is the subject of Chapter 16. For now, it suffices to recognise that gravity is not a quantum field awaiting discovery, but the structural description of the geometry that emerges whenever Phase 1 is resolved into Phase 2.

The Early Galaxy Formation Problem (#14)

The unusually early, massive galaxies revealed by JWST present a serious problem for Λ CDM because, within that framework, there simply isn't enough time for such systems to form. The problem only arises if the Λ CDM timeline is treated as a literal record of events, stretching cleanly from a hot dense beginning to the present day. If that whole history is taken for granted, the new data look impossible.

2PC removes that expectation. If the early universe is not an earlier physical era but an observational surface reconstructed from the present cosmic state, then this question, like so many others, changes shape. Instead of asking how real galaxies managed to assemble so quickly in a young universe, we must ask why the structures that appear on the reconstructed early surface have the mass and maturity that they do.

One potential explanation is that the pattern we see is a by-product of the same selection constraints that shape the rest of the embodied branch. If complex life requires a particular set of large-scale conditions, and those conditions depend on the way structure develops over cosmic time, then only those retrodicted early surfaces that support a viable long-term path to embodiment will appear within Phase 2. This does not mean that LUCAS needed distant galaxies to form early in a literal past. It means that in the version of the cosmos where everything was just right for LUCAS to evolve in the Milky Way, conditions across the entire cosmos favoured the early development of this kind of galaxies. Indeed, here might be the answer to the question we were left with in Chapter 11 about which "cosmic egg" is selected for realisation. There are infinite possible cosmoses in the Pythagorean ensemble, but the existence of these galaxies suggests that the selected cosmos (the next in the queue to hatch) is the one where conscious life evolves in the shortest period of time. The selected cosmos is not the one that "wants" consciousness fastest, but the one whose global structure reaches representational stability earliest under the same physical laws. This would mean that not only do we live in a cosmos where everything is just right for the evolution of conscious organisms, but the one where everything is so exceptionally just right that the evolution

of LUCAS also happens as quickly as is physically possible. If future observations revealed that galaxy formation was delayed until significantly later cosmic times (say, $z < 5$ for massive galaxies) this would present a challenge to the 'fast track' interpretation, suggesting that selection favors other constraints over temporal efficiency. This is a second novel empirical prediction.

Ton 618

A similar pattern appears in the case of extreme quasars such as Ton 618, whose central black hole is so massive that, under standard assumptions, it appears to violate the Eddington limit (the maximum rate that objects can acquire mass). Within Λ CDM this again presents itself as a timing and growth problem: there is not enough time for a black hole to reach such a mass through ordinary, radiation-limited accretion without invoking exotic seeds or prolonged super-Eddington phases. As with the JWST galaxies, the difficulty only arises if the inferred early universe is treated as a literal record of gradual physical construction.

Readers will be familiar with the logic now. Ton 618 does not require a historical episode in which physical limits were repeatedly broken; it requires only that the retrodicted early surface of the embodied branch contains black holes whose integrated mass and luminosity histories are globally consistent with the present state of the cosmos. The Eddington limit constrains local accretion dynamics, not the structure of a globally selected spacetime history reconstructed from an already-realised world. Seen this way, extreme early quasars are not anomalies but further expressions of the same selection constraint suggested by the JWST data.

The Black Hole Information Paradox (#16)

In standard physics, the Black Hole Information Paradox arises because Hawking radiation appears thermal and featureless, suggesting that information falling into a black hole is permanently lost, in violation of quantum mechanical unitarity. In 2PC, this paradox dissolves once we recognise that black holes mark not a breakdown of physics, but the boundary beyond which Phase-2 realisation cannot be sustained.

Phase 1 and Phase 2 Perspectives

From the perspective of Phase 1, no information is ever lost. The global superposition preserves complete unitarity; all degrees of freedom, including those that appear to fall into black holes, remain encoded in the mathematical structure of possibility space. The paradox arises only if we assume that Phase 2 constitutes the entirety of reality.

Phase 2, by contrast, is local and selective. It contains only those structures that have become entangled with collapse loci – regions where conscious systems have instantiated determinate states. The interior of a black hole, beyond the event horizon, is therefore a Phase-2 exclusion region: not because physical laws fail there, but because the conditions required for sustained representational coherence cannot ultimately be maintained. Information entering the horizon does not vanish; it simply remains in Phase 1 as uninstantiated formal structure.

Hawking Radiation and the Page Curve

Hawking radiation emerges in standard theory as a quantum effect associated with the horizon. In 2PC, this radiation exists formally within Phase 1, while only the portion that becomes entangled with existing collapse loci is realised in Phase 2. Crucially, the radiation realised in Phase 2 carries only information about exterior degrees of freedom that were already part of the realised manifold. It does not encode interior microstates, because those states were never instantiated in Phase 2 to begin with. This resolves the tension with the Page-curve requirement of unitarity. From the Phase-1 perspective, unitarity is perfectly preserved: all information remains encoded in the timeless

superposition. From the Phase 2 perspective, radiation appears thermal and information appears lost, reflecting the locality limits of realisation rather than any breakdown of quantum mechanics. Hawking radiation is thus the observable signature of the boundary between realised and unrealised domains.

The Infalling Observer

The status of an infalling observer requires careful clarification. In 2PC, the event horizon is not a physical barrier and does not violate the equivalence principle. A freely falling observer experiences no local discontinuity at the horizon, exactly as general relativity predicts. The horizon is instead an ontological boundary defined by global causal structure, not by local physics. Phase 2 participation requires sustained representational coherence within an open network of bidirectional causal exchange. Crossing the horizon does not abruptly terminate this participation; rather, it initiates an irreversible process. Beyond the horizon, the observer's causal embedding progressively shrinks. Increasing portions of the realised universe become permanently inaccessible, and the predictive self-model can no longer maintain global coherence with the wider collapse network. As this causal isolation deepens, the storm of micro-collapses sustaining Phase 2 existence becomes progressively destabilised. From the internal perspective, therefore, there is no sharp event or instantaneous "firewall." Instead, Phase 2 participation fades asymptotically as the system loses the ability to sustain coherent self-modelling within the realised manifold. Consciousness terminates as the conditions required for stable collapse can no longer be met. The observer's informational structure is not destroyed; it persists within Phase 1 as unrealised formal possibility.

Conceptual Significance

Black holes in 2PC are neither destroyers of information nor sites requiring exotic quantum-gravity mechanisms to preserve unitarity. They are regions where the Phase-2 manifold reaches a causal limit beyond which sustained representational coherence becomes impossible. The information paradox disappears once we distinguish two domains: Phase 1, where all information is eternally preserved within the mathematical ensemble, and Phase 2, where only that which can be coherently instantiated becomes real. The event horizon marks the global boundary of Phase-2 participation, while the dissolution of realised structures occurs gradually as causal isolation erodes the conditions required for ongoing collapse. Hawking radiation is the thermal trace of this boundary, carrying away the limited information that was ever part of the realised manifold, while the full informational content remains intact within the timeless vault of possibility.

The Fermi Paradox (#15)

The Fermi paradox has always carried a strange mix of curiosity and dread. The universe is old and wide and full of places where life should have taken hold, and yet every radio dish comes up empty. Under 2PC this becomes painfully clear. The primordial wave function could collapse only once, because collapse is not a repeating physical glitch but a metaphysical resolution of contradiction. When that resolution happened, the universe stopped behaving like a Phase 1 generator of branching possibilities and settled into the single embodied history that could host a coherent subject. Phase 1 had something like the power of an unbounded search, but that power only existed before consciousness appeared. Once a subject emerged and valuations had to be unified, that search space was no longer available. The cosmos we inhabit is the one that made it through the threshold. It is not physically impossible that it could happen again, but it is so extraordinarily improbable that we can rule it out (this is the very same "almost certainly false" that appears in the subtitle of *Mind and Cosmos*). So the question of where everyone is becomes almost rhetorical. If any separate locus of conscious existence ever did arise, it would almost certainly belong to a

metaphysically disconnected branch that we can never encounter from inside this collapse. In the world we can interact with, there is no one else. *We are it.*

This is the third novel empirical prediction. If we were to find strong evidence for a second example of life having evolved, anywhere in the cosmos, we would have legitimate reasons for doubting the two phase model, and if we find conscious life then 2PC is almost certainly false.

Neogeocentrism

It turns out the old geocentrists weren't as far off as we like to think. Ever since the shift from Ptolemy to Copernicus, people have assumed that once Earth was removed from the centre of the solar system it must also have lost any special place in the larger cosmos. Later we decided the universe didn't have a centre at all, and most people accepted that even if they never quite understood what it meant. It felt harmless to believe the universe was everywhere and nowhere at once, so no one worried about it. 2PC changes this in a very simple way. If consciousness emerged only once, and if that emergence is what brought Phase 2 into being, then the point of origin for the embodied cosmos is not a spatial coordinate at all but a locus of subjectivity. Our planet does not sit at some geometric centre, but it is the only location that carries the structures capable of embodiment. Earth is the *metaphysical* centre of the cosmos.

The fate of Λ CDM

From the perspective of the old paradigm, which is also the perspective of academic cosmology, the claims I have made in this chapter are outrageous. I am expecting them to be ignored, dismissed and ridiculed, in that order. But the reality is that if we change the metaphysical foundation – if we reject physicalism and replace it with the hierarchical neutral monism I have described in this book, then the problems that Λ CDM cannot solve either disappear, or are fundamentally reframed. The truth is that Λ CDM is like a ship holed beneath the waterline, and as fast as cosmologists try to seal the holes, more holes appear. This ship is heading for the bottom. The only thing that has been keeping it afloat is the absence of an empirically adequate paradigm to replace it. Or rather I should say *was*, because The Two-Phase Cosmology is that paradigm. This will be no ordinary paradigm shift though, because physicalism is going down with Λ CDM, because Λ CDM is the best model that physicalism can support.

Chapter 15: A new kind of paradigm shift

As a civilisation and as a species, we find ourselves in very deep trouble indeed. I have already written a book about this, and one is enough. To be more specific, *The Real Paths to Ecocivilisation* was about the intersection between those problems and the more Western-specific problems that I explored in Part One of this book, especially the deep fragmentation of knowledge and its many debilitating consequences. We aren't just in existential trouble; we are also ideologically stuck. Some sort of "cultural reset" is required.

Unfortunately, as I made clear in RPE, it is too late to do this the "nice way": a lot of what is coming is going to be both involuntary and apocalyptically bad. However, a central point I was making in that book was that this cannot be a reason to stop thinking – to stop caring about our choices. I hope that the connection between this and the arguments in the present book should now be obvious. When I wrote RPE I was still lacking the threshold mechanism, so I didn't make any strong connection between that threshold and the metaphysics of decision-making. I was making the simpler point that there are still a vast range of possibilities – that the future is not written, and the difference between the worst possible outcome and the best possible outcome is enormous. I argued that while most of our physical, practical problems are guaranteed to get worse (and probably much worse), there's very real justification for a new sort of optimism with respect to the way our societies operate: it is entirely possible that we're at the very lowest point in a cultural cycle: the point where everything seems completely hopeless, which also happens to be the moment just before light finally appears at the end of a long tunnel.

A very brief history of Western worldviews

Western history goes something like this:

The deepest roots of Western civilisation can be found in ancient Greece and Rome. The Greeks invented philosophy, politics and fine art, and though they were great experimenters in civilisation-building, they never scaled it up beyond the city state. The Romans invented the republic, perfected the art of expansionism and sorted out much of the “nuts and bolts” of large-scale civilisation, but their version of civilisation was, even by our own unimpressive standards, deficient in terms of morality and genuine spirituality. Brutality, cruelty and slavery were considered normal, while mercifulness was regarded as a sign of weakness. Then along came Christianity. The details of exactly how and why this happened have become historically obscured by the mythology of Christian origins. Christians generally regard the mythology as history, while non-Christians tend towards the idea that mythology is all there is: that Jesus may not even have existed. What is not in doubt is that the Romans tried but failed to suppress the new movement, and as the Empire stagnated and decayed Christianity became the “new attractor”. When Rome finally fell to the barbarians, Europe entered a “dark age”, the church hoarded power, and the philosophies of the ancients were either forgotten or subsumed into the grand theological synthesis of Augustine and Aquinas. While the ancients emphasised rational inquiry even at the expense of moral and spiritual concerns, the medieval world (at least in theory) placed morality and spirituality at the centre – which required the subordination of reason to theological authority. Nonetheless, Western civilisation had for the first and only time arrived at a common foundational worldview.

The next great revolution was arguably triggered by the socio-economic fallout from the

population crash of the Black Death, but is more often considered to have begun with the Renaissance: the rediscovery of important lost works of ancient philosophy, mostly in the form of translations made by Islamic scholars, and the re-ignition of fine art. This ultimately led to the Scientific Revolution and the Enlightenment – the mature fruit of the Renaissance conviction that moderns could surpass the ancients. This was also the time that capitalism began to replace feudalism as a socio-economic system, and when representative democracy began to replace absolute monarchy. It was the birth of the modern Western world – and of the globalised civilisation we currently know (even though that now includes most of the world, not just the West). The common foundation had failed, and there was now a growing number of incompatible and mutually contradictory worldviews. A monumental battle raged between materialistic science and the fractured remains of Christianity. Modern civilisation brought with it many wonderful things. Our world has been transformed in many positive ways – it hasn't all been problems. But certain things have also gone horribly wrong, and there has been a major philosophical and political response to these failures. That response is called Postmodernism.

Postmodernism

Postmodernism emerged in the twentieth century not as a unified philosophy, but as a broad cultural reaction against the assumptions, aspirations, and blind spots of modernism. To understand its importance in the present context, we must appreciate that this was not mere academic contrarianism, but a necessary reckoning with some of the unintended consequences of modern thought. Modernism, as a philosophical and cultural project, placed its faith in reason, science, universal truth, and progress. It assumed that history had a direction, that knowledge could be built on secure foundations, and that the human condition could be improved indefinitely through technological advancement and rational governance. The Enlightenment had promised emancipation from superstition and tyranny through science and reason, and modernism was its cultural heir. Postmodernism rejected this optimism, finding within it the seeds of domination and exclusion. Thinkers like Michel Foucault, Jean-François Lyotard, Jacques Derrida and others relentlessly attacked the very idea of “universal truth”, arguing that so-called universal values often mask the interests of particular groups (i.e. white, male, heterosexual elites). The Enlightenment promise of reason, they argued, had been co-opted by institutions of power: science had become instrumentalised, rationality bureaucratized, and knowledge weaponised in service of empire, industry, and the state.

Lyotard's famous definition of postmodernism is “incredulity toward metanarratives”: postmodernism is deeply skeptical of modernism's grand stories about progress, freedom, or objective truth, claiming that these narratives excluded, suppressed, and silenced other ways of knowing. Reason and science were not considered to be neutral arbiters of truth; they were situated, contingent, and interwoven with systems of power. This critique was especially powerful when applied to the ecological crisis. From a postmodern perspective, modernity's faith in control and mastery over nature was itself the root of our environmental problems. The modern subject – autonomous, rational, and separate from the world – conceived nature as a standing reserve of resources to be exploited. Postmodernism helped expose this anthropocentric delusion, and pointed instead toward indigenous, feminist, and other "marginalised epistemologies" that had long emphasised relationality, reciprocity, and respect for the more-than-human world.

Postmodernism did not merely deconstruct modernism's assumptions; it also intentionally disrupted its language. Derrida's analysis of texts convinced many people that meaning is never fixed – that words always carry within them the possibility of contradiction, ambiguity, and slippage. The stable categories and clear boundaries of modernist thought were recast as illusions. Language, and by extension knowledge itself, was declared to be a kind of game: contingent, contextual, and open-ended. This refusal to offer new certainties in place of the old ones has led directly to relativism and nihilism, but postmodernists will argue that this criticism misses the point:

they will say that postmodernism was not a doctrine of despair, but an ethic of humility which has shown us that no system of thought is above critique, and that pluralism, diversity, and dialogue are better foundations for living together than unquestionable monolithic truths or rigid hierarchies. In the context of this broader historical arc, postmodernism must be understood as an immune response to modernity's overreach. It cleared the ground, exposed the rot, and made space for something genuinely new to emerge. Unfortunately, **it did not build anything new**, and does not provide us with any of the tools we need to begin that task. Postmodernism was the point where the old epistemic frame lost contact with any possible grounding.

As things stand, the various forms of "post-postmodernism" are not a new paradigm but a symptom of paradigm exhaustion. They are, at best, a patchwork of tentative proposals driven more by the yearning for coherence than by its discovery. If you want to explore post-postmodernism, then some internet searches that can throw some light are GameB, Metamodernism, Second Renaissance, Unified Theory of Knowledge (UTOK), the Liminal Web, Polycrisis, Metacrisis, Meaning Crisis, Sensemaking communities, Integral Theory and post-progressivism. However, if I had to recommend a starting point it would be the work of Iain McGilchrist. That is a long list, and there are all sorts of ideas involved, and some of them contradict others (so they cannot all be right). What all of these movements, thinkers, and frameworks are searching for is a new integrative worldview that can move beyond both modernism and postmodernism. This new worldview must reconnect meaning, science, spirituality, and systems thinking in order to navigate civilisational crisis and co-create a sustainable, coherent, and life-affirming future.

The old paradigm

The Real Paths to Ecocivilisation was another attempt to make progress towards the new paradigm, and it did so in a way that none of these people or movements have done – not even McGilchrist. I suspect he would view my ideas as overly conceptual and systematising, whereas I view his vision as a profound diagnosis of modernity's epistemic pathology which is nevertheless vulnerable to vagueness and romanticism because it stops short of proposing a rigorous alternative framework. While the historical progression of ancient->medieval->modern->postmodern is undeniable, I am rejecting the idea of this as an intellectual or personal progression. This rejection is also what happens if you apply postmodern thinking to itself, but I have no intention of doing that; I am not interested in declaring postmodernism to be self-refuting and trying to slide backwards into modernism, for that would just ping me back towards postmodernism and set up a metamodern "oscillation". Instead, the post-postmodernism I'm advocating rejects the historical narrative described above. I don't see postmodernism as a paradigmatic improvement on modernism. Rather, I see postmodernism as the ultimate expression of the failure of the old paradigm: the vanguard of the West's descent into collective nihilistic psychosis. In McGilchristian terms it is what happens when the brain's left hemisphere finally loses the plot completely and starts attacking its own flawed creations, but with no attempt to recover the meaning, context and coherence that only the neglected right hemisphere can provide. This is where the Two-Phase Cosmology comes in, because it finally allows us to make sense of quantum mechanics (and it has taken exactly a century). To understand how this can be the key to fixing Western thinking, we must return to the moment when our epistemological troubles really began: David Hume's *A Treatise of Human Nature: Being an Attempt to Introduce the Experimental Method of Reasoning into Moral Subjects*, which was first published in 1739.

During the Renaissance, reality had been divided into mind and matter first by Galileo, and then by Descartes. This was done intentionally: both of these revolutionary thinkers saw the material world as that which can be quantified and measured, and mind as the realm of that which cannot. It was the business of the new "natural philosophy" to investigate the material world, in an attempt to reduce it to the workings of natural laws, and the method of investigation involved the systematic elimination of everything subjective – indeed, *that was the whole point*. In the decades

that followed, philosophers searched for a way to put the rest of philosophy on as firm foundations as those of materialistic science, but rather than finding an agreement about how this could be done, a long battle was fought between the defenders of two conflicting approaches to grounding that system. The empiricists argued that knowledge must start with observations of reality, the rationalists argued that it must start with pure reason, and Hume was the first person to put his finger on the epicentre of the entire problematic. His goal in the *Treatise* was to provide solid foundations for a science of mind ("moral subject" here means a conscious human), but even though his analysis was miles ahead of anybody else at that time, he ran into a logical problem that totally defeated him. He felt he had compelling reasons for believing two contradictory things. The first was that for all we know, we could be brains in vats: how could we ever transcend "the veil of perception" and know anything about a world that lies beyond? The second was that in order to be able to experience an external world (as we evidently do), then it must be the case that objects in the external world have a causal effect on our minds – there has to be a direct chain of cause and effect from an external object to the subjective experience of that object. The contradiction appears to be very real: either there are mind-external objects which can causally penetrate the veil of perception, or there aren't, right?

Hume never found a solution to this problem. His conclusion to that section of the *Treatise* is one of the most tortured passages ever written. After all his extensive and faultless reasoning, he could give no justification for believing anything positive at all. He could do no better than say that the world strongly appears to be the way it appears, which in terms of epistemology is scant improvement on Descartes' argument that we should believe what our senses are telling us because God wouldn't deceive us. What we suppose to be the real world "enlivens some ideas beyond others" – our perception of an external world just "feels stronger" than merely internal mental activity or dreams. Without this feeling, we'd have no reason to reject solipsism or subjective idealism, but feelings are "so inconstant and fallacious" that this sort of principle will surely lead us into errors. Such a methodology is never going to be scientific, that is for sure, and yet it is only feelings, experience and habit which makes us "reason from cause and effect". In other words, it is only because we are so familiar with the world behaving as if causality is real that we believe in it "and 'tis the same principle, which convinces us of the continu'd existence of external objects, when absent from the senses." The problem is that belief in the reality of causally effective mind-external objects is "natural and necessary in the human mind". How could we function without it? But how can this be reconciled with the reasoning which forces us towards skepticism?

"How then shall we adjust those principles together? Which of them shall we prefer? Or in case we prefer neither of them, but successively assent to both, as is usual among philosophers, with what confidence can we afterwards usurp that glorious title, when we thus knowingly embrace a manifest contradiction?"

It was exactly this contradiction that prompted Immanuel Kant to write the *Critique of Pure Reason* (1781/87)(CPR), in which he made the paradigm-defining move of dividing reality into phenomenon (reality as it appears to us) and noumenon (reality as it is in itself) instead of mind and matter. Kant argued that science can only tell us about phenomena, and that noumena were forever not just unknowable but uncognisable. For Kant, space and time are conditions for human experience – they are the frame for physical phenomena, and we have no reason to believe they exist in noumenal reality. Kant's masterpiece provided the foundation for modern Western philosophy, but this also marked the point where it began to split into two divergent "traditions". One branch led via Schopenhauer and Nietzsche to what is now called "Continental philosophy" (which fully embraces contradiction), and the other led to "analytic philosophy" (which attempts to be rational but still hasn't found an acceptable solution to Hume's problem).

It is essential to understand the context in which Hume and Kant were working. This

was the golden age of materialistic science. Newton's *Principia* had blown the old ways of thinking to smithereens and both of them were trying to bring the subjective world of consciousness, and therefore the whole of reality, onto a similarly secure footing. Absolutely nobody had the slightest inkling that one day we would discover that Newtonian physics is not the "final" description of reality after all, but without that piece of information even the genius and precision of Hume and Kant stood no chance of identifying the correct solution to the philosophical problems of their day.

Hume, Kant and quantum mechanics

Now let us imagine that history had played out differently. Let's imagine that physics had advanced at a much more rapid pace and that in the time between Hume's *Treatise* and Kant starting work on the CPR, quantum theory had been discovered. Now, instead of having to find a way to solve Hume's problems in the light of an apparently undeniable fact that Newtonian physics is the one true description of physical reality, let us imagine Kant was aware of the Measurement Problem. In order to rid science of the "quantum leap" from superposition to a single state, von Neumann had been forced by logic to propose a conscious observer outside of the physical/quantum system. But in the real history, Hume and Kant were dealing with a physical model which was a direct match for the phenomenal world of "ordinary" material objects. That was what set the problem up: what could be more obviously correct than to have a physical theory which describes the reality we actually experience? And yet we now know this assumption is wrong. In the imaginary history the situation is very different – here science provides a physical model which does not describe the phenomenal-material world. Instead of being a classical realm of material objects, it is the non-local realm of the wave-function. If you think about it this way, then Hume's problem disappears. We can now map physics onto reality with no difficulty at all (although ironically we can do this only if we're willing to abandon physicalism in favour of neutral monism). We can simply say that the unobserved world – the "real world" which is out there "beyond the veil of perception" is the world described by the equations of our best physics: the uncollapsed wave function. In 2PC we call this "Phase 1". Only when reality collapses into Phase 2 does the world of ordinary material objects appear. Hume's veil lifts not because the wavefunction is the noumenon but because the noumenal/phenomenal split is reconfigured once we recognise that superposition is the metaphysical substrate from which the single actual world is continuously resolved. What Kant called "noumenon" is therefore shown not be so completely unknowable after all. There are some very specific things which we do not know about it, but 21st century technology is built on our knowledge of Phase 1.

The new paradigm

The point of this extended detour into the history of Western philosophy is this: **if 2PC is right, then we can re-ground realism.** This might sound implausible given that I'm saying that before the Cambrian Explosion, reality as we understand it didn't even exist, but why should that matter in the present context? The nature of the reality of the cosmos 600 million years ago isn't what matters if we're talking about regrounding realism for civilisation today. All that matters is that we can find a new way to agree that the reality we actually find ourselves in is real, and not an illusion or a social construction.

This clears a new pathway out of the thicket. Not just the beginning of a new search because we've concluded that we must move beyond postmodernism and don't know what comes next, but more like a final correction and completion of modernism. Postmodern anti-realism isn't a stage which both people and societies need to pass through on the way to some strange promised land where modernism and postmodernism perpetually undermine each other. Instead we need to accept that the anti-realistic relativism of postmodernism was based on a gigantic but entirely unavoidable mistake. In fact, even before the discovery of quantum mechanics, science was always telling us things about an objective world, beyond the veil of perception. Hume's instinct was right: of course

we are aware of real objects. What he got wrong, because he had no chance of getting it right, was that instead of being like the Phase 2 material objects we directly perceive, they exist in a way that is non-spatiotemporal, and consists of all physical possibilities at the same time.

This provides new scope to for a collective agreement that reality is real after all, even if it is not the normal material world that modernism assumed it to be. It means that science is not just another perspective, as laden with power dynamics as political and religious ideologies. It does this not in the reductive manner of materialism, but in a way which affirms the reality of consciousness and everything that comes with it – especially meaning and value. A new kind of Western thinking can be founded on the idea that Cartesian substance dualism and its Kantian solidification into knowable phenomena and unknowable noumena were both **epistemological mistakes**. Cartesian dualism was a hopeless oversimplification, and Kant was trying to solve a problem without the empirical knowledge necessary to solve it. The real dualism is between possibility and actuality, and this dualism is underwritten by the Void. While the original insight came from Eastern philosophies, this is something new, and something very much Western: a post-postmodern neo-Kantian non-panpsychist neutral monism.

As for the people and movements I listed earlier in this chapter, I believe some will prove to be closer to the mark than others. Beware of attempts to dress the old paradigm up in new clothes. These come in two main forms. The first I've already discussed – the metamodern attempt to smuggle (in plain sight) postmodern anti-realism into the new paradigm by framing it as one pole in an “oscillation.” Make no mistake: if you mix realism and anti-realism (or modernism and postmodernism), the result will be postmodern anti-realism. One times minus one equals minus one. Metamodernism is the bargaining stage of grief for bereaved postmodernists.

The second form is exemplified by UTOK, which from the perspective of 2PC is the corresponding stage for bereaved materialists: an old-paradigm-style naturalism that remains fundamentally hostile to genuine spirituality, offering a psychological control system in place of metaphysical insight. If we are to find a true path to ecocivilisation, we must begin by understanding where we went wrong – not only politically or economically, but philosophically. *The Real Paths to Ecocivilisation* was an attempt to lay the philosophical foundation for a civilisation that can truly endure – because it is grounded not in illusion, but in a reawakened relationship with reality. Endlessly oscillating between the failures and limitations of modernism and the dead marshes of postmodernism is not going to cut the mustard.

Chapter 16: The praeternatural and the New Epistemic Deal

I must now specify a new set of definitions. The old distinction between natural and supernatural does not work in a cosmos where the laws of physics permit two different kinds of causality. The "praeternatural" and Phase 1 of 2PC is what Wolfgang Pauli searched for and couldn't find. He called them "prae-physics" and "the Unus Mundus", respectively.

Naturalism (or "metaphysical naturalism") is a metaphysical view that everything happening in reality can be reduced to (or explained in terms of) the laws of nature (including laws we are yet to discover). Naturalism is logically entailed by physicalism (all physicalists are naturalists) but the reverse is not true (some naturalists aren't physicalists). If the physical world is the only thing that exists, then there is no theoretical or conceptual space for anything else that could affect it (an agent of free will, for example, or God) so naturalism is logically entailed. Naturalism differs from determinism in that it can accommodate objective randomness – the future does not have to be fully determined, but anything not determined must be random (*really* random in every case, not just apparently so or only in some cases).

Some people use the terms "laws of physics" and "laws of nature" interchangeably, but they have slightly different meanings. The meaning of "laws of physics" is straightforward enough. "Laws of nature" is broader and encompasses all natural phenomena, including not just physical laws, but also biological, chemical, and maybe even some aspects of social or psychological phenomena that are considered "natural" in a broad sense. The laws of nature can include principles like natural selection in biology or chemical reaction rates in chemistry, in addition to the laws of physics.

Given the above definition of naturalism, one might think it is straightforward to define supernaturalism as an opposing term: the belief that naturalism is false – the belief in forms of causality (or something resembling causality) that can't be reduced to laws of nature; the belief that *something else* is going on. Relatively mild examples include free will, synchronicity and karma. More extreme examples are Young Earth Creationism and most of the miracles in the gospels, such as the Resurrection and the Feeding of the 5,000.

Not so long ago this definition would have been perfectly adequate. For the whole period between Newton and Einstein (inclusive) physics made absolute predictions about future observations. The laws of classical physics and relativity leave no wiggle room for anything else (which is why Einstein declared that God does not play dice with the universe). In a world where those are our best theories of physics then supernaturalism necessarily involves a breach of laws – it has to involve something that intervenes in the clockwork causality of the material world to produce effects that are inconsistent with the laws of physics. This would be true of everything listed above as supernatural – the mild examples are as incompatible with classical-relativistic physics as the extreme ones.

However, modern physics is quantum physics, and quantum physics differs from what went before in several important ways. One of these is that the predictions made by quantum physics are probabilistic rather than absolute. This element of probability provides enough wiggle room to render the straightforward definition of supernaturalism useless.

MWI in particular helps to demonstrate the implications of the probabilistic nature of quantum theory. Even though most timelines in an MWI multiverse are “normal”, there will always be a small minority where all manner of strange things happen, including not just occasional improbable events but what appear to be co-ordinated sequences of them. The reason this matters is that the alleged supernatural phenomena on the “mild” list are consistent with the laws of physics: they require nothing more than events, or sequences of events, which are extraordinarily improbable. Take synchronicity for example. Carl Gustav Jung defined synchronicity as “an acausal connecting principle”. He equated it with the Tao, and connected it to Taoism’s divinatory system of the I Ching. Synchronicity manifests as combinations of events that are both exceptionally improbable and meaningful in some way, particularly but not always associated with personal psychological or spiritual development. Jung’s best known example occurred during a therapy session with a rational and highly educated woman who had been resisting dealing with her emotions. She was telling him about a dream in which she had been given a piece of jewellery in the form of an Egyptian golden scarab beetle – a symbol of rebirth and transformation. Jung heard a tapping on the closed window behind him and saw a flying insect knocking against the window from outside. He opened the window and caught the creature, which turned out to be a beetle virtually identical to the golden scarab, rarely found at that latitude, especially trying to get into dark rooms during the day. The improbability of such an event occurring at precisely that time allowed the woman to stop intellectualising matters and from that moment she began to make progress.

Science can give us no reason to believe in any such thing, because it gives us no way to distinguish between synchronicity and objectively random chance (only a subjective value judgement is appropriate here). But neither does science give us any reason to rule it out as impossible. It can’t do, because MWI is consistent with science, and if it is true then there are timelines where events that fit the description of synchronicity happen to everybody, all the time. In fact in this case it would not actually be synchronicity, because there would be no need to posit an acausal connecting principle as the explanation. The point is that such events are physically possible even if MWI is false.

Now contrast this with the things on the “extreme list”. There are no MWI timelines where YEC happened and none where it is possible to adequately feed 5,000 people with five loaves and two fishes. Even in the unimaginably vast MWI multiverse there are no events which breach the laws of physics.

Because of the level of uncertainty regarding the interpretations of quantum theory, the definition of supernaturalism given above is not adequate. The phenomena on the extreme list still fit the old definition (they require a suspension of the laws of physics) but the phenomena on the mild list no longer belong in that category. On the one hand they do not contradict the laws of physics, but on the other they are not explained by those laws or reducible to them. There is a conceptual space and need here for two categories where currently there is only one.

The first we might call **contra-scientific supernaturalism**. These are forms of causality that require an outright suspension of the laws of nature/physics. A person who truly believes that 5,000 people can be adequately fed with two fishes and five loaves cannot be expected to accept the epistemic authority of science. This presents serious challenges when attempting to incorporate mainstream Christianity – or any other worldview that involves belief in contra-scientific supernatural phenomena – into an epistemic framework fit for the future. If this type of supernaturalism were possible, why wouldn’t God fix our world right now? If you can feed 5,000 people like that, then you can feed eight billion and more. Why doesn’t He conjure up another planet for us? The second we might call **probabilistic supernaturalism**: forms of causality involving exceptionally improbable events that are consistent with the laws of physics but that are not explained by them or reducible to them.

Paranormal

The term **paranormal** is sometimes considered a subset of the supernatural. “Supernatural” is an older term and historically associated with religion, whereas “paranormal” refers to alleged oddities that are beyond current scientific understanding but might not always be. Perhaps we could use it to refer to probabilistic supernaturalism, but I fear that might encourage people to attempt to use scientific methods to investigate, prove or disprove these probabilistic phenomena, which I believe to be a misleading waste of time. If the phenomena in question exist at all, and they are fully amenable to scientific investigation, then why didn’t incontrovertible evidence emerge long ago? What has actually happened is an endless dispute about how to interpret borderline results which seem clear enough to people who already believe in such things but not clear enough to convince the skeptics. The believers accuse the skeptics of repeatedly raising the bar to ensure that all positive evidence is rejected as inadequate. The skeptics say the bar *needs* to be high because if you are making extra-ordinary claims then you need extra-ordinary evidence (the “Sagan standard”).

I doubt this situation is going to change any time soon. I think the way forwards is for people to agree to disagree (at least for now), accept that the alleged empirical evidence is too borderline to convince the skeptics, and give up trying to resolve the unresolvable. If probabilistic supernatural phenomena are real then they involve causality (or something that resembles it) which is fundamentally different to anything science has ever successfully demystified. Jung didn’t even try to invoke synchronicity (in his narrower sense) under laboratory conditions, and to suggest such a thing is possible is to misunderstand the nature of the alleged phenomena. By definition, these are things that happen to people for a reason, and the reason is never that they are taking part in a scientific experiment designed to prove or disprove the existence of the phenomena themselves.

I have no good use for the term “paranormal”.

Praeternatural and hypernatural

There is another term available, which might be more appropriate than “probabilistic supernaturalism”. St Thomas Aquinas (c1225-1274) was the greatest of Catholic philosophers, and from his time onwards he was considered the only philosopher to have got Catholic Christianity “correct”. Aquinas claimed that God sometimes works miracles, but nobody else can – that magic is possible, with the help of demons, but is not properly miraculous. This distinction can strike modern people as odd, but from the time of St Thomas until the 16th century people had a different set of causal categories to us. We think of magic and miracles as synonymous – both are supernatural as opposed to natural causality. They had three categories instead of two – “supernatural” and “miracle” were terms reserved for acts of God involving a suspension of the natural order. Magic was categorised as praeternatural, which means “beyond nature”. Even though demons were involved, this was technically a manipulation of the natural order rather than its outright suspension. Praeternatural phenomena *could* have been entirely the result of natural causality, but aren’t. Magic – aka witchcraft or sorcery – was considered very real and most evil, hence this period is well known for the widespread persecution of alleged witches (of both sexes, but more frequently women). By the mid-18th century the term “praeternatural” had fallen out of use, and it eventually gained a modern non-metaphysical meaning of “so talented it’s spooky”.

I don’t like the term “contra-scientific supernatural” – it is too cumbersome. “Probabilistic supernatural” isn’t quite right either. “Probabilistic” is fine, but anything “supernatural” sounds like it involves a suspension of the laws of physics. I therefore use “hypernatural” rather than “contra-scientific supernatural” and “praeternatural” for “probabilistic supernatural”. The term “supernatural” thereby disappears with the old paradigm, which will make very clear in any context whether we are talking about the old concept of supernatural, or the new concepts I am suggesting should replace it.

From here onwards:

Naturalism is belief in a causal order in which everything that happens can be reduced to (or explained in terms of) the laws of nature.

Hypernaturalism is belief in a causal order in which there are events or processes that require a suspension or breach of the laws of nature.

Praeternaturalism is belief in a causal order in which there are no events that require a suspension or breach of the laws of nature, but there can be exceptionally improbable events that aren't reducible to those laws, and aren't random either. Praeternatural phenomena *could* have been entirely the result of natural causality, but aren't.

A New Epistemic Deal

In order to start making real progress towards fixing our world's problems, we need to change the way we think and communicate. The New Epistemic Deal is a proposed agreement to facilitate serious discussions about ecology, politics, economics, morality, spirituality and countless other topics without the discussion being derailed by people who refuse to start with the known facts about reality. We must do this without scientistically ruling out things that aren't supported by science or reason but aren't in conflict with them either.

1: Ecocivilisation is our shared destiny and guiding goal.

Ecocivilisation represents a vision of a society that harmonises human activity with ecological principles. This is not a utopian ideal but a necessity dictated by the realities of ecosystems and evolution. The claim that ecocivilisation is our destiny is pre-political, transcending specific ideologies or systems. The social, political, and economic structures of ecocivilisation are not part of this definition, but the core premise is clear: civilisation must work *ecologically* to endure.

2: Consciousness is real.

Consciousness – our individual interface with reality – is the one thing each of us can be absolutely certain exists. It is through consciousness that we perceive existence and recognise that anything exists at all. As such, consciousness must serve as the starting point for exploring what exists beyond our subjective experience and for discerning the boundaries of what we know and what we don't.

3: Epistemic structural realism is true.

Scientific knowledge tends towards truth. We acknowledge that there is such a thing as an objective reality, external to human minds, about which science provides structural knowledge that is reliable, albeit with certain important qualifications. We reject the idea that all scientific knowledge is merely provisional, or as subjective as non-scientific forms of knowledge. We affirm the epistemic privilege of science.

4: Both materialism and physicalism should be rejected.

Materialism cannot account for consciousness. Physicalism either suffers from the same problem, or it implies things that most physicalists reject, in which case it is not much use as a piece of terminology. Both materialism and physicalism restrict our models of reality in such a way that they are never going to be able to satisfactorily account for everything we have justification for believing exists.

Principles five and six are more convoluted, but their justification should now be clear:

5: The existence of praeternatural phenomena is consistent with science and reason, but apart from the cases of psychogenesis and free will, there is no scientific or rational justification for believing in it/them either. The only possible justification for belief is subjective lived experience. Psychogenesis and free will are structural parts of the system – there's no point in being skeptical about them. Everybody is aware of their own consciousness and felt metaphysical freedom – we do not consider this “spooky”. The same does not apply to synchronicity or any other alleged praeternatural phenomena.

6: We cannot expect people to believe things (*any things*) based solely on *other people's* subjective lived experiences. There will always be skeptics about any alleged praeternatural phenomena and their right to skepticism must be respected.

Principles seven and eight are closely related, but sufficiently distinct to warrant the inclusion of both.

7: There can be no morality if we deny reality.

If there actually is an objective reality, and we can actually know things about it, then if we start our moral reasoning with anything other than reality we are engaged in fake morality – we will be arguing about what would be morally right and wrong in some ideal reality rather than the real one that we have to figure out how to share. And if the people we are having moral disagreements with are actually dealing with reality, while we are not, then they are engaged with real morality and we are claiming moral high ground we have no right to claim.

8: Science, including ecology, must take epistemic privilege over economics, politics and everything else that purports to be about objective reality.

Principle seven is specifically about morality. Principle eight is about everything that matters – it is about practical reasoning as well as moral reasoning. It demands that the whole of science, including the whole of ecology, the limits to growth and the reality of ecological overshoot, must be acknowledged before serious discussion starts about *anything at all*. It should be considered *immoral* to come to any negotiating table demanding concessions from others before you are willing to accept reality. Growth-based economics and politics are dangerous nonsense, and for anybody who understands that, engaging with them while failing to persistently challenge their false assumptions is an immoral act.

The mechanism of the praeternatural: Competition Resolved Collapse (CRC)

Technical Definition of Competition Resolved Collapse

After the Embodiment Threshold has been passed, collapse unfolds as a storm of micro-collapses inside the subject's specious present. This storm is the ongoing process that sustains embodiment.

Each micro-collapse (**ci**) is a local stabilisation of entangled alternatives. The hazard rate (which explains **when** the world collapses (driven by attention and value)) for a micro-collapse is:

$\lambda_i(t) = \lambda_0 * [1 + \alpha_V * V_i(t) + \alpha_P * P_i(t) + \alpha_A * A_i(t) + \alpha_C * C_i(t)]$

$$\lambda_i(t) = \lambda_0 [1 + \alpha_V V_i(t) + \alpha_P P_i(t) + \alpha_A A_i(t) + \alpha_C C_i(t)]$$

Competing micro-collapses share overlapping support in Hilbert space, and the realised one minimises the inconsistency functional (which path the world takes (the one that causes the least "logical pain")). :

$$F[c_i] = |\langle \Psi | \hat{O}_{c_i} | \Psi \rangle - \hat{V}_{c_i}|^2 + \beta D(\rho_{SE} \| \rho_S \otimes \rho_E)$$

$$\mathcal{F}[c_i] = |\langle \Psi | \hat{O}_{c_i} | \Psi \rangle - \hat{V}_{c_i}|^2 + \beta D(\rho_{SE} \| \rho_S \otimes \rho_E)$$

The dynamics across the specious present follow a rate-modulated stochastic field, defined by the master equation:

$$\dot{\rho}_S = \mathcal{L}_U[\rho_S] + i \sum (\lambda_i(t) * (M^i * \rho_S * M^{i\dagger} - 0.5 * \{M^{i\dagger} * M^i, \rho_S\}))$$

$$\dot{\rho}_S = \mathcal{L}_U[\rho_S] + i \sum \lambda_i(t) \left(M^i \rho_S M^{i\dagger} - \frac{1}{2} \{M^{i\dagger} M^i, \rho_S\} \right)$$

This equation represents the structural completion of the model: it shows how the standard laws of physics (\mathcal{L}_U) and the subject's internal valuations (λ_i) combine to create a single, unfolding reality. This "master equation" is the final rule that combines it all into a moving picture of reality.

How Competition-Resolved Collapse closes the multi-agent hole

Once the Embodiment Threshold has been crossed, collapse is no longer a single global event. It becomes an ongoing process: a dense, overlapping storm of local stabilisations occurring within the specious present of embodied agents. Each micro-collapse is a resolution of nearby physical alternatives under valuation, prediction, attention, and coherence constraints. This process sustains the continuity of experience and action for a subject.

The problem that CRC is designed to address arises only when more than one such subject exists. If conscious agents are genuinely real, and if their valuations matter, then situations inevitably arise in which two or more agents assign incompatible values to the same unfolding physical situation. One agent intends one outcome; another intends a mutually exclusive outcome. If collapse were driven independently by each agent's valuations alone, the theory would permit incompatible realised worlds – private realities with no principled mechanism forcing agreement. That is the multi-agent hole. CRC as defined here is deliberately minimal. It exists to close this structural inconsistency, *not to explain meaning*.

CRC closes this hole without introducing a new force, a coordinating intelligence, or a hidden global selector. The key move is simple but non-negotiable: when agents interact, their micro-collapse processes become entangled. Their valuation-weighted collapse hazards no longer operate on disjoint physical supports. Instead, overlapping regions of Hilbert space are jointly constrained by all participating agents' valuations and predictive structures. At that point, incompatible continuations are not merely in conflict at the level of desire or belief. They are in conflict at the level of representability. A continuation in which both incompatible valuations are simultaneously realised cannot remain dynamically coherent, because it would require the same shared physical degrees of freedom to stabilise in mutually exclusive ways.

CRC resolves this by treating collapse as a competition among overlapping micro-collapses. Each candidate stabilisation carries a cost, measured by an inconsistency functional. This functional penalises two things: mismatch between the physical outcome and the valuations applied to it, and breakdown of coherence between the agents and their shared environment. The realised micro-collapse is the one that minimises this combined inconsistency. Importantly, this is not a vote, a negotiation, or a moral arbitration. No agent “wins” because it is stronger, more numerous, or more important. The outcome is whichever continuation can be jointly stabilised by the entangled system of agents and environment with the least representational contradiction. The others simply fail to stabilise and are not realised. This is how a single shared world is maintained. Agents remain free to value, intend, and predict independently. Conflicts are not prevented. But when those conflicts concern the same physical degrees of freedom, only outcomes that can be jointly embodied survive the collapse process. The shared world is not imposed from outside; it is the residue of what multiple embodied perspectives can coherently sustain together.

CRC therefore does exactly one thing: it ensures that multi-agent reality does not fragment. It does not explain why particular symbolic meanings arise, why certain patterns feel significant, or why some coincidences strike us as meaningful rather than accidental. Those phenomena occur after the closure CRC provides, within the space of stabilised shared reality. In other words, CRC closes the door on solipsistic branching, but it does not seal the building. What follows, once a shared world is secured, is a further question: how symbolic systems, shared myths, and authorised interpretive frameworks modulate which micro-collapses are preferentially stabilised within that world.

The Problem of Meaning and Value (#30)

This book ends with the solution to the final problem of the original thirty, and the completion of the life's work of the the sharpest thinker, and the deepest philosopher, of all of the founders of quantum mechanics.

Pauli's prae-physics and the Unus Mundus

Wolfgang Pauli was widely recognised as the "conscience" of physics, and he spent much of his intellectual life circling a question he never managed to bring to ground. It was not a technical problem, nor a failure of imagination, but a boundary he could see clearly and could not cross without giving up the very standards of intelligibility that had made him what he was in the first place. What he sought was not a hidden variable or a missing equation. By the time he entered his long correspondence with Jung, Pauli had already accepted that the quantum world had broken the classical picture beyond repair. What troubled him was even deeper. Physics described lawful regularities with extraordinary precision, yet the form of those laws, and the very fact that the world came in meaningful structures at all, seemed to lie beyond physics itself. At the same time, psychic phenomena (dreams, symbols, synchronicities) exhibited order without causation, and significance without mechanism. Pauli suspected that both domains pointed toward something prior to the split between matter and mind: a pre-physical ordering domain in which meaning and law had not yet separated.

This was his "prae-physics" (a term he only used in personal letters, never in official academic work). It was not intended to be mystical in the popular sense (Pauli was deeply suspicious of anything that smelled of occultism or animism). He wanted rigour, symmetry, and formal constraint. He imagined that beneath both quantum law and psychic structure there must exist a deeper level of reality, which could be expressed in a language not yet invented but still continuous with physics. He believed that if such a level existed, it would show itself through structural correspondences, shared symmetries, or a common mathematical grammar binding psyche and matter together. From the standpoint of 2PC, this is exactly where Pauli's search became

impossible. He never relinquished the assumption that whatever grounds reality must itself be structurally articulable. Even when he spoke of complementarity between mind and matter, he assumed that the complementarity itself belonged to a unified order that could, in principle, be described. He was willing to enlarge physics, but not to fracture ontology. He believed there had to be continuity all the way down.

Phase 1, as it is defined in 2PC, violates that assumption. It is not a deeper layer of law or an abstract structure awaiting discovery. It is not governed by symmetry, causation, probability, or meaning. It is timeless, non-causal, non-valuational, and non-instantiated. There is no privileged way to decompose it, no preferred basis, and no fact of the matter about which distinctions are significant. It does not contain laws from which other laws descend. It contains only the space of what is logically and physically consistent, without commitment. Pauli was searching for the laws behind the laws, but Phase 1 doesn't have any. This is why he repeatedly encountered what felt like a void and then recoiled. From within Phase 2 reasoning, Phase 1 does not look like a domain at all. It looks like indeterminacy, or nothingness, or mysticism stripped of content. Pauli explored each of these interpretations and rejected them, because none of them could be reconciled with his demand for intelligibility.

The decisive move that Pauli never made, and arguably could not make, concerns collapse. Pauli treated collapse as something that must ultimately be explained by a principle – perhaps a deeper symmetry, perhaps a psycho-physical correspondence, perhaps an extension of complementarity. He assumed that if collapse were real, there must be a reason for it that could be stated as part of the world's order. But in 2PC, collapse is not governed by a deeper law. It is not a causal or stochastic process, and not a symmetry-breaking event in the usual sense. It is a meta-logical resolution forced by incoherence: when incompatible valuations arise within a unified representational perspective, reality cannot remain uncommitted. Something must give, and what gives is superposition itself.

Pauli searched for a principle behind collapse, but 2PC identifies collapse as what happens when the demand for coherence can no longer be deferred, and that difference is decisive. The same limitation shaped Pauli's struggle with synchronicity. He was convinced that synchronistic events were not mere illusion, yet he could not locate them anywhere within physics. Without a distinction between global, pre-instantiated possibility and local, post-collapse reality, synchronicity could only be interpreted in two ways: either as purely a projection of the human mind, or as evidence of a hidden cosmic order arranging events acausally. Pauli rejected both options. The first felt dishonest; the second felt like a return to pre-scientific teleology. Within 2PC, the resolution is straightforward. Meaning does not exist before collapse. It exists only after collapse, as a constraint on how local micro-collapses stabilise within Phase 2. Synchronicities (and the dreams to which Pauli attached so much significance) are not messages from a deeper order; they are patterns that emerge when an agent's symbolic commitments shape the coherence of their local collapse dynamics. Pauli lacked the ontological place to put this, because he tried to locate meaning prior to instantiation. He reached the boundary of what his conceptual commitments allowed him to say and there he stood, fully aware of the abyss, and choosing not to step into what would have looked, from his standpoint, like nonsense.

The semantic constraint: Extended Competition Resolved Collapse (E-CRC)

In the spirit of Principle 6 of the New Epistemic Deal, this is what I choose to believe about praeternatural phenomena beyond psychogenesis and free will. I am not asking anybody else to believe it, because it is, in part, based on my own experiences. Standard CRC is sufficient to complete my model, but I believe it also leaves something important out. I must emphasise that what follows is my best guess. I leave open the possibility that it may turn out to be too

conservative/skeptical, though I doubt there will ever be any objective evidence to support anything stronger.

Extended CRC: Where Meaning Emerges

While CRC handles the physical and representational structure, humans live in worlds rich with symbols, myths, and beliefs. Here is where Extended CRC (E-CRC) enters. E-CRC is not a new law or force; it is a semantic constraint layer. It shapes which micro-collapse outcomes are allowed to stabilise, based on the agent's authorised symbolic frameworks. Crucially, symbolic systems (whether religious, divinatory, or cultural) do not have to be true. Belief is *permission*, not assent. In other words, this is another level of constraint, consistent with the whole of the rest of this model of reality. If the agent permits a symbolic framework, outcomes consistent with it are allowed; incompatible outcomes fail to stabilise. Over the storm, this filtering produces the phenomenology of meaningful coincidences, divinatory answers, and culturally specific experiences. No micro-collapse alone carries meaning. Meaning emerges only as a pattern across many micro-collapses, stabilised by the symbolic code the agent has authorised. Nothing is sent or transmitted, nothing exists beyond the agent's experience, and no archetype or cosmic intelligence is invoked.

Structured randomness (as in tarot, the I Ching, or other divinatory systems) interacts naturally with E-CRC. A random draw partitions possibilities; the symbolic framework selects which interpretations can stabilise. For a believer, the system appears to "answer." For a skeptic, the same draw produces no coherent pattern. This is not superstition; it is the local, structured resolution of representational ambiguity under symbolic constraints. Religious and anomalous experiences follow the same principle. The form of experience reflects the symbolic authorisation of the agent. A Christian experiences Christian-coded phenomena, a Hindu experiences Hindu-coded phenomena, a UFO believer experiences alien-coded phenomena. These experiences are fully real for the agent, but they do not generalise beyond their local domain.

Synchronicity Reinterpreted

Jung's notion of synchronicity (meaningful coincidences between inner states and external events) is recovered without metaphysical excess. A synchronicity is a temporally local pattern of micro-collapse resolutions. The inner state and the external event are both part of the storm; when they stabilise coherently under an authorised symbolic framework, the agent experiences coincidence. No causal link is forged. No external message is sent. The "meaning" is a property of the pattern, not of the world itself.

Pauli searched for a deeper, pre-physical principle to connect mind and matter. Jung reified archetypes and postulated acausal ordering. Both were blocked by assumptions: Pauli assumed structural representation in Phase 1; Jung assumed that meaning must exist outside the subject. Two-Phase Cosmology shows why both were partially right and partially blocked. There is a "phase cut": Phase 1 is timeless, non-causal, and non-instantiated; Phase 2 supports consciousness and local collapse. Meaning arises after collapse, constrained by symbols authorised by the agent. There is no bridge, no hidden law, and no external ordering principle, only the storm resolving under local constraints. The symbolic imagery of Pauli's dreams, which he described in great detail to Jung, was derived from physics. For example, "radioactivity" symbolised the process of inner transformation, and "spin" symbolised the rhythm of the universe. The origin of this meaning was not a "psychoid" realm where meaning and matter were one, but Pauli's own belief system.

This architecture explicitly prevents inflation of ontology. There are no mind-independent archetypes, no (Bohmian) active information acting on matter, no cosmic intelligence, no objective divinatory truth, and no cross-agent enforcement of symbols. Attempting to introduce these would break the formalism as I specify it.

Formalising Extended CRC (E-CRC)

To bridge the gap between abstract collapse and lived meaning, we need an Extended Competition-Resolved Collapse (E-CRC). This does not change the dynamics of physics; it introduces semantic constraint-based filters.

The Core Principles of E-CRC

Agent-Relative Constraint: Collapse outcomes are constrained by an agent's authorised symbolic frameworks. These frameworks are not "truths" but permitted interpretive codes.

Semantic Authorisation: A symbolic system (e.g., tarot, mythic schemas) influences collapse only if the agent grants it "permission" by treating it as meaningful.

Structured Randomisation: Meaning emerges by coupling a stochastic substrate with a pre-structured semantic space. E-CRC does not bias randomness; it suppresses outcomes that are incoherent within the authorised code.

No Ontological Export: Outcomes produced under symbolic constraint are strictly local to the agent. They provide no evidence of external supernatural entities or shared metaphysical structures.

Formal Scaffold

To represent this mathematically, we define the coherence functional (C_A) which measures representational incompatibility.

Definitions:

Omega: Set of physically admissible Phase 2 continuations.

VA: Agent's valuation structure.

SA: Agent's authorised symbolic system.

CA: Coherence functional (measures incoherence/incompatibility).

The Resolution Rule:

$\omega^* = \operatorname{argmin}_{\omega \in \Omega} [C_A(\omega \mid V_A, S_A)]$

$$\omega^* \in \operatorname{argmin}_{\omega \in \Omega} C_A(\omega \mid V_A, S_A)$$

The Micro-Collapse Storm and the Embodiment Threshold

The Embodiment Threshold (ET) is the transition that instantiates the subject. It is pre-symbolic and pre-belief. It occurs because valuation contradictions render unitary evolution logically incoherent. Post-ET, consciousness is a temporally extended storm of local micro-collapses. Each micro-collapse resolves a local representational ambiguity to sustain the self-model. E-CRC governs this resolution.

$\omega_t = \operatorname{argmin}_{\omega \in \Omega_t} [C_A(\omega \mid V_A, S_A)]$

$$\omega_t \in \operatorname{argmin}_{\omega \in \Omega_t} C_A(\omega \mid V_A, S_A)$$

Coda

Two things fill the mind with ever new and increasing admiration and awe, the more often and steadily we reflect upon them: the starry heavens above me and the moral law within me.

I no longer need to choose between them.