How to wrap your head around the most mind-bending theories of reality

From the many worlds interpretation to panpsychism, theories of reality often sound absurd. Here's how you can figure out which ones to take seriously

ARE there vastly many near-duplicates of you reading vastly many near-duplicates of this article in vastly many parallel universes? Is consciousness a fundamental property of all

matter? Could reality be a computer simulation? Reader, I can hear your groans from here in California.

We are inclined to reject ideas like these on the grounds that they sound preposterous. And yet some of the world's leading scientists and philosophers advocate for them. Why? And how should you, assuming you aren't an expert, react to these sorts of hypotheses?

When we confront fundamental questions about the nature of reality, things quickly get weird. As a philosopher specialising in metaphysics, I submit that weirdness is inevitable, and that something radically bizarre will turn out to be true.

Read more

Stephen Hawking's final theorem turns time and causality inside out

Which isn't to say that every odd hypothesis is created equal. On the contrary, some weird possibilities are worth taking more seriously than others. Positing Zorg the Destroyer, hidden at the galactic core and pulling on protons with invisible strings, would rightly be laughed away as an explanation for anything. But we can mindfully evaluate the various preposterous-seeming ideas that deserve serious consideration, even in the absence of straightforward

empirical tests.

The key is to become comfortable weighing competing implausibilities, something that we can all try – so long as we don't expect to all arrive at the same conclusions.

Let us start by clarifying that we are talking here about questions monstrously large and formidable: the foundations of reality and the basis of our understanding of those foundations. What is the underlying structure of the universe and how does human consciousness fit into it?

It isn't that we don't have theories that attempt to answer these questions. Rather, we have too many, and no straightforward way to put them to the test – nor any realistic prospect of doing so any time soon. This is where weirdness and bewilderment reign.

To be more specific, my argument is that all broad-ranging attempts to articulate the fundamental structure of reality, no matter how soberly we approach them, inevitably become both bizarre and dubious. Bizarre, in that they defy common sense. And dubious, in that they allow for reasonable doubt. These twin theses of "universal bizarreness" and "universal dubiety" in fundamental matters of metaphysics and cosmology form the basis of my latest book, The Weirdness of the World.

Quantum mechanics

Consider interpretations of quantum mechanics. We have a set of equations that describes the behaviour of the subatomic particles that comprise matter and the forces through which they interact. That behaviour is famously strange. But the equations of quantum mechanics aren't particularly dubious, having passed every experimental test we have thrown at them.

It is when we attempt to assess what they tell us about the nature of reality that we confront their full bizarreness and dubiety. According to some interpretations, for example, the act of observing a quantum object collapses its "wave function", the mathematical entity that encodes the probabilities of all possible outcomes. In these interpretations, particles exist in a quantum fog until our observations distil them into hard-edged classical reality.

The thought experiment known as Schrödinger's cat illustrates the bizarreness of this view. A cat is trapped in a box with a poison that may or may not have been released. It exists in a "superposition" of alive and dead until the moment someone peers into the box and collapses the wave function.

According to "objective collapse" interpretations, resolving the wave function into one outcome needn't rely on human

observation and is itself a matter of chance. Schrödinger's cat could, in principle – though it is unlikely – remain in the alive-dead superposition even after we observe it. Still other interpretations, known as Bohmian mechanics after the physicist David Bohm, posit a faster-than-light "pilot wave" that determines outcomes in quantum systems.

Rethinking reality: Is the entire universe a single quantum object?

In the face of new evidence, physicists are starting to view the cosmos not as made up of disparate layers, but as a quantum whole linked by entanglement

Then there are interpretations that predict no collapse at all.

According to the "many worlds" interpretation, the cat remains always in the superposition of alive and dead. When we open the lid, we join the superposition, and hence from that moment also exist in two possible states: a version of us sees the cat as alive and another sees the cat as dead, and these two versions of us go on to lead different lives in different "worlds" – one in which the cat died, with all the consequences of that, and one in which the cat lived.

Such world-splitting would be extremely common, occurring, for example, almost every time a photon passes through polarised sunglasses. Hence, according to many worlds, there are vastly many near-duplicates of you reading vastly many near-duplicates of this article in vastly many parallel universes.

That might seem utterly preposterous. But the same could be said of all other interpretations of quantum mechanics, and that is the point. They are all bizarre and dubious, in different respects and for different reasons. Many worlds's bizarreness can't, therefore, be decisive evidence against it. No non-bizarre interpretation is currently available, nor likely to be waiting in the wings.

That shouldn't be surprising. Common sense evolved to help us negotiate the world as we see it – comprising middle-sized objects and low velocities – and we ought not to expect it to help much beyond that domain.

How certain are you that we don’t live in a computer simulation?

Saha Entertainment/Getty Images

Nor do any of the interpretations entirely compel belief. For the foreseeable future, there is no decisive empirical test we can run that will straightforwardly settle the question of which is the correct one.

So, something ragingly weird must be true about the quantum realm. But which of the various preposterous-seeming possibilities is in fact correct? The truth is that we are in no good position to know. We have reached the limits of science and philosophy, leaving us all – experts or otherwise – in a similar position.

But here's the thing: some theories do fit a bit more naturally with the empirical evidence, and each has its own virtues, such as simplicity, elegance and fruitfulness in suggesting new research directions. It is on these grounds that we can compare bizarre theories and pick favourites.

Many worlds, for example, is more sharply at odds with common sense than most of the other interpretations of quantum mechanics. Since we have to start our thinking somewhere, that is a strike against it. Violations of common sense have to be earned.

On the other hand, many worlds has advantages in terms of theoretical elegance. There is no wave function collapse to worry about, so no need for a model of what really happens during collapse. And, arguably, it fits with the empirical evidence more naturally than other ideas, since there are no speculative collapse dynamics that remain undiscovered.

Did time flow in two directions from the big bang, making two futures?

Why time only flows forwards is one of the great mysteries of physics. A new idea suggests that it actually went in two ways from the big bang – and, even more radically, that time emerges not from entropy, but from the growth of structure

In contrast to many worlds, with objective collapse interpretations there is always a chance of collapse, too rare to be seen in experiments on individual particles but common enough to be nearly instant in macroscopic systems. Other interpretations attribute collapse to the act of observation – but what is “observation”, and why should it

cause collapse?

Although bizarre, such ideas aren't quite as bizarre as many worlds. On the other hand, the speculative claims about collapse make them less empirically simple and elegant. However, there is a way in which collapse theories are simpler, since they prevent the proliferation of worlds. They are, in that respect, leaner as far as metaphysics is concerned.

How you weigh up such competing virtues is down to you. Some of us are happy to let common sense go if it yields a cleaner idea. Others more stubbornly hold to it until forced to do otherwise. Scientists and non-scientists can reasonably disagree about such matters, within certain parameters.

Those parameters are important. The idea that Zorg the Destroyer causes quantum collapse from his hiding place in the galactic core, for instance, has no merits. You should roll your eyes at that. But when it comes to many worlds, you can appreciate the general structure of competing considerations and reach your own educated guess.

Panpsychism

We can apply similar thinking to panpsychism, the idea that consciousness is a fundamental property of matter, always present, rather than something new that arises only when

matter is combined in the right way.

The people who give credence to panpsychism do so for two main reasons. First, it isn't at all obvious how something as distinctive as full-colour conscious experience could arise from humdrum physical matter, and panpsychism offers a solution. If all matter has some small degree of consciousness, it is easier to see how it might combine into the more complex consciousness we know.

There are other ways to explain consciousness, of course. They include substance dualism, which holds that mind and matter are fundamentally distinct; a priori physicalism, which argues that consciousness somehow emerges from ordinary physical matter; and illusionism, the idea that the very concept of consciousness is unscientific. But each faces its own troubles.

As far as panpsychism goes, the second reason to give it credence lies in the challenge of drawing the line between organisms or systems with and without consciousness. Logically, there are four options. Everything is conscious, as in panpsychism; nothing is conscious; there is a hard line between conscious and non-conscious systems; or there is a fuzzy line between conscious and non-conscious systems. Any strikes against the final three options should make us favour panpsychism.

Presumably, you know that you are conscious. That rules out the nothing-is-conscious option. Arguably, it is theoretically inelegant and empirically untenable to posit a hard line between conscious and non-conscious organisms, given that living things exist on a continuum of psychological capacities and behaviours. That rules out the hard-line option. And it is difficult to understand how there could be a fuzzy boundary between conscious and non-conscious systems. What would it even mean for a system to be “kind of” conscious, standing somehow between having subjective experiences and not having subjective experiences? Ruling out those three options would leave you with panpsychism.

Quantum weirdness isn't weird – if we accept objects don't exist

We can grasp the truth about the quantum world, says physicist Carlo Rovelli – as long as we abandon our most cherished assumption about what's real and what's not

But we can reasonably resist the panpsychist conclusion. In my own research, I have defended the existence of a fuzzy line between conscious and non-conscious systems. We ought to accept – though it is difficult and maybe even impossible to conceive – that between no experience and some little flicker of experience lies a half-something that is neither a nothing nor a something.

I have phrased that paradoxically to bring out its poor fit with our conceptions of consciousness. As soon as we imagine an entity as conscious, we imagine it as having some consciousness or other, however vague or minimal, or we don't imagine it as conscious at all. I can see why people might want to resist the fuzzy-line view.

So how do we choose? Here, again, we are past the point where knowledge and expertise can settle the question, and we can all exercise our independent judgement within a range of reasonable options.

Panpsychism elegantly handles both the puzzle of how

consciousness fits into a physical world and the puzzle of how to draw the boundary between conscious and non-conscious entities. Much like many worlds, it strains mightily against common sense, while having an attractive simplicity.

In the quantum multiverse, there are many versions of you

TIMOTHY A. CLARY/AFP via Getty Images

Which leaves the last of the three questions we began with: how can we assess the idea that we are all living in a computer simulation?

Clearly, we are still far from a consensus scientific theory concerning the nature of consciousness. But many potential explanations suggest that it is, at least in principle, possible

to create consciousness in artificial intelligence systems. So, suppose you allow that there is a good possibility, in principle, of such a consciousness arising somewhere, someday, using some technology. And if you allow that, you should probably also allow that some of these conscious AI systems live entirely within simulated environments, receiving inputs from computational outputs – roughly akin to what a human experiences when they don virtual reality goggles. Such AI systems would be living inside simulations, though they might not know it.

Simulation theory

If we grant this pair of possibilities, then there is a basis for wondering whether many such AI systems exist – and we can wonder if we are among them.

There is also a basis to resist this, of course. Maybe no AI system could ever be conscious. Maybe genuinely conscious AI systems living in simulated realities would be expensive, and technologically advanced designers would have little motive to create them. Maybe the fact that we haven't noticed any glitches in our reality constitutes some empirical evidence against the simulation hypothesis.

Alternatively, maybe conscious AI systems living in virtual realities would be relatively inexpensive for sufficiently advanced societies. Designers might create them for

science, entertainment or because they regard it as ethically worthwhile to create worlds full of rich, meaningful life.

Stacking up the possibilities, we might regard it as highly unlikely that we live in a simulation. But it isn't Zorg territory. The possibilities aren't so remote and absurd that we should disregard them. Perhaps we ought to be 99.9 per cent confident we don't live in a simulation. But in light of our ignorance about consciousness and fundamental cosmology, and in light of the bizarreness and dubiety of all theories in this area, it seems hard to justify a level of 99.999999 per cent confidence that the simulation hypothesis is false.

At which point, you might be wondering if we are resigned to confusion. Will we never know the right interpretation of quantum mechanics, or the grounds of consciousness, or whether we live in a simulation?

Not necessarily. Science can reveal answers to questions that previously seemed unsolvable. What once appeared bizarre can become comfortable and familiar. But as we await new empirical evidence to support or refute many worlds or panpsychism, we can delight in our ignorance. We can celebrate the diversity of wild options and the weirdness of the world. How boring it would be, after all, if there were no bizarre metaphysical possibilities to wonder about.

Eric Schwitzgebel is a philosopher at the University of California, Riverside. His latest book is The Weirdness of the World