New Scientist

WEEKLY 3 August 2019

ANTI-AGEING PILL
Can new supplement make you healthier for longer?

VIKING SECRETS
Genetics reveals true origin of invaders

PEAK DINOSAUR
Why the golden age of discovery has only just begun

REALITY: THE GREATEST ILLUSION OF ALL

How evolution has blinded us to the truth about the world

WATER FROM AIR
Tech that could quench the planet's thirst

PLUS ARCTIC WILDFIRES / **NEW MARMOSET** / BIONIC EYE / **IS IT BETTER TO BUY LOCAL?** / BLOWING UP BRAIN CELLS

No3241 £4.95 CAN$7.99

Do we see reality?

The idea that our perception matches reality is intuitive – but it may be blinding us to deeper truths, says cognitive scientist **Donald Hoffman**

LIFE insurance is a bet on objective reality – a bet that something exists, even if I cease to. This bet seems quite safe to most of us. Life insurance is, accordingly, a lucrative business.

While we are alive and paying premiums, our conscious experiences constitute a different kind of reality, a subjective reality. My experience of a pounding migraine is certainly real to me, but it wouldn't exist if I didn't. My visual experience of a red cherry fades to an experience of grey when I shut my eyes. Objective reality, I presume, doesn't likewise fade to grey.

What is the relationship between the world out there and my internal experience of it – between objective and subjective reality? If I'm sober, and don't suspect a prank, I'm inclined to believe that when I see a cherry, there is a real cherry whose shape and colour match my experience, and which continues to exist when I look away.

This assumption is central to how we think about ourselves and the world. But is it valid? Experiments my collaborators and I have performed to test the form of sensory perception that evolution has given us suggest a startling conclusion: it isn't. It leads to a crazy-sounding conclusion, that we may all be gripped by a collective delusion about the nature of the material world. If that is correct, it could have ramifications across the breadth of science – from how consciousness arises to the nature of quantum weirdness to the shape of a future "theory of everything". Reality may never seem the same again.

The idea that what we perceive might differ from objective reality dates back millennia. Ancient Greek philosopher Plato proposed that we are like prisoners shackled in a fire-lit cave. The action of reality is happening out of sight behind us, and we see only a flickering shadow of it projected onto the cave wall.

Modern science largely abandoned such speculation. For centuries, we have made stunning progress by assuming that physical objects, and the space and time in which they move, are objectively real. This assumption underlies scientific theories from Newtonian mechanics to Albert Einstein's relativity to Charles Darwin's theory of evolution by natural selection.

Natural selection, you might think, gives a simple reason why our senses must get it largely right about objective reality. Those of our predecessors who saw more accurately were more successful at performing essential tasks necessary for survival, such as feeding, fighting, fleeing and mating. They were more likely to pass on their genes, which coded for more accurate perceptions. Evolution will naturally select for senses that give us a truer view of the world. As the evolutionary theorist Robert Trivers puts it: "Our sense organs have evolved to give us a marvellously detailed and accurate view of the outside world."

"For centuries we have made stunning progress by assuming things are real"

The truth of such statements can be tested with mathematical rigour using the tools of evolutionary game theory, introduced in the 1970s by John Maynard Smith. In this theory, different strategies for coping with the natural world can be set against each other in simulations to see which approaches are fitter – in the sense of producing more offspring.

In the case of perception, we can study how "truth" strategies, which see objective reality as it is, fare against "pay-off" strategies, which see only survival value. Take oxygen. Too much or too little oxygen in the air kills us; a narrow range keeps us alive. Now imagine living in an environment where the level of oxygen varies considerably, and you have to make survival judgements about whether to venture outside.

For the sake of this example, the amount of oxygen in the air is taken to be an objective truth. You might imagine seeing it on a colour scale from red for low to green for high. That's the truth strategy: you know the truth, but you don't know if you'll die. A pay-off strategy would mean seeing red colours for levels of oxygen that kill you, and green for those that don't. You see what you need to survive, but don't see the objective truth of how much oxygen there is.

MARINA LOEB

The objective truth I started seeing a decade ago, in simulations conducted together with my graduate students Justin Mark and Brian Marion at the University of California, Irvine, is that evolution ruthlessly selects against truth strategies and for pay-off strategies. An organism that sees objective reality is always less fit than an organism of equal complexity that sees fitness pay-offs. Seeing objective reality will make you extinct.

If this seems hard to swallow, suppose you are writing a novel on a laptop, and the novel's icon on the desktop is green, rectangular and in the centre of the screen. Does this mean that the novel itself is green, rectangular and in the centre of your laptop?

Of course not. The desktop interface is there to mask a complex reality of software, circuits and digital 1s and 0s to provide a simple way to interact with it. If you actually had to flip computer bits to write a novel, you would switch to pen and paper.

Reality is virtual

That, evolutionary game theory predicts, is what evolution has done for us. Natural selection has given us sensory systems that are a simplifying user interface for the complexity of the world. Space, as we perceive it around us, is a 3D computer desktop, with tables, chairs, the moon and mountains icons within it.

In other words, our senses constitute a virtual reality. If you play the video game *Grand Theft Auto* with a virtual-reality add-on, you see a 3D world with 3D objects, such as a black steering wheel in front of you. If you turn your head, however, the steering wheel disappears. Indeed, it ceases to exist, because it only exists when we are looking where it should be in the simulation. The reality that exists – circuits and software again – is utterly unlike a steering wheel. But it prompts you to create a steering wheel when it is needed, and to destroy it when it isn't.

In like manner, we create an apple when we look, and destroy it when we look away. Something exists when we don't look, but it isn't an apple, and is probably nothing like an apple. The human perception of an apple is a data structure that indicates something edible (a fitness pay-off) and how to eat it. We create these data structures with a glance, and erase them with a blink. Physical objects, and indeed the space and time they exist in, are evolution's way of presenting fitness pay-offs in a compact and usable form.

But hang on, drop the apple. A lion on the African savannah isn't just an icon in your

Our senses tell us only what we need to know to survive

interface. It has agency, and can kill you, so it must be objectively real.

I wouldn't mess with a lion, for the same reason I wouldn't carelessly drag the green icon of my novel to the virtual recycle bin. Not because I take that icon literally, and think the novel is green and rectangular. But I do take that icon seriously: if I drag it to the bin, I could lose all my work.

The objection that a lion must be objectively real because anyone who looks over there sees a lion that we can all agree looks like a lion – so it isn't unique to our subjective experience – isn't a valid one, either. Humans agree about what we see because we have all evolved a similar interface. The interfaces of some other species, such as prey mammals, may have icons for lions that are similar to ours, and that guide actions similar to ours, such as keeping far away from them.

That leaves the fact that treating our observed, subjective reality as objective reality has allowed us to create scientific theories – frameworks that allow us to make predictions about how the world works, and so are presumably part of an objective reality that exists outside our heads. But here too there are hints from deep within science itself that perception and reality don't match.

Quantum theory is our best physical theory of fundamental reality. But with its counter-intuitive effects of "spooky action at a distance" and the perennial mystery of the dead-yet-alive Schrödinger's cat, it drives a coach and horses through cherished ideas from our classical realm of experience: that objects have definite values of the properties pertaining to them, that those properties don't depend on how they are observed, and that influences propagate no faster than light.

This is jolting if we assume that objects and their measurable properties are part of an objective reality. But it is no surprise if we think of objects and their properties as data structures created as needed to represent fitness pay-offs. In this case, the values of properties will depend on when and how we create them.

This approach aligns with the quantum-Bayesian interpretation of quantum theory, or QBism, in which the uncertainty inherent in quantum observations is all in the minds of the observers. As three pioneers of QBism, Christopher Fuchs, David Mermin and Rüdiger Schack have put it, "A measurement does not, as the term unfortunately suggests, reveal a pre-existing state of affairs. It is an action on the world by an agent that results in the creation of an outcome – a new experience for that agent. 'Intervention' might be a better term."

PLAINPICTURE/RUDI SEBASTIAN

If our team's evolutionary ideas are true, that might lend momentum to models of quantum theory that see quantum states, and the mathematical and interpretational structures around them, as "epistemic" – reflecting not necessarily reality, but just our state of knowledge of it.

But it goes further. Even perceptions as seemingly fundamental as space and time might not actually be part of objective reality. That insight could inform our search for theories that unite the two great theories at the heart of modern physics.

For decades, we have tried and failed to reconcile quantum theory with general relativity, Einstein's theory of gravity that

"Even perceptions as fundamental as space and time might not be part of objective reality"

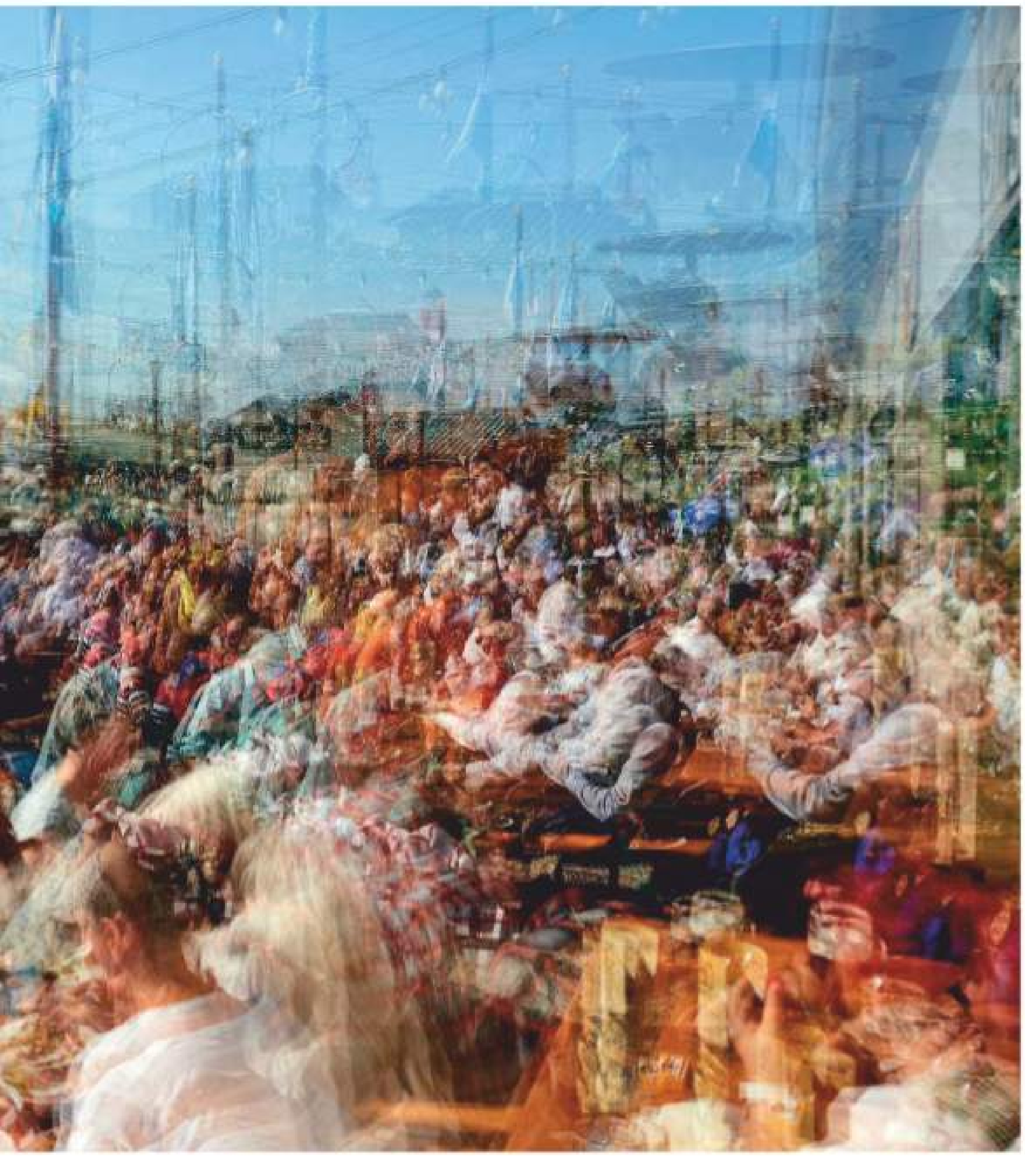

Donald Hoffman is a professor of cognitive sciences at the University of California, Irvine. His research interests span visual perception, artificial intelligence, evolutionary psychology and the hard problem of consciousness. His book on human perception, *The Case Against Reality*, is published this month.

dictates how the universe works on large scales. At a very basic level, these theories fail to agree on the nature of space and time.

General relativity demands that space-time, the four-dimensional structure that space and time together form, is smooth and continuous, whereas a quantum description requires a pixelated description. As the theoretical physicist Nima Arkani-Hamed has said: "Almost all of us believe that space-time doesn't exist, that space-time is doomed, and has to be replaced by some more primitive building blocks." Admittedly, no one yet knows what those might be – but our insights suggest the hunch they must be replaced is right.

It isn't just in physics where we may need to overhaul our ideas about reality to make progress. Another is in solving the "hard problem" of consciousness. This problem of how and why our brains generate conscious experience remains intractable despite centuries of thought. As biologist Thomas Huxley put it in 1869: "How it is that anything so remarkable as a state of consciousness comes about as a result of irritating nervous tissue, is just as unaccountable as the appearance of the djinn, when Aladdin rubbed his lamp."

The brain-exciting technology of transcranial magnetic stimulation (TMS) illustrates how little progress we have made. Suppose we place a TMS unit near your scalp, on the right side of your head, near an area of the occipital cortex called V4. We turn on the device, and its strong and focused magnetic fields inhibit neural activity nearby. All colour drains away from the left half of your visual world; you see only shades of grey. We turn off the device, and the colour seeps back in.

Chocolate and vanilla

Neuroscience has turned up hundreds of such correlations between patterns of neural activity and specific conscious experiences. Most attempts to explain these correlations assume that the neural activity causes, or somehow gives rise to them. But how, precisely? What neural activity causes the taste of vanilla, and why doesn't it cause the taste of chocolate? In a network of interacting neurons, how exactly do changes in voltage, or in the flow of sodium, potassium and calcium ions through pores in neural membranes, create an individual conscious experience?

There are no theories, and few plausible ideas. But if we are trying to find the answer to the problem of conscious experience in the firing of neurons in space and time, when those neurons themselves are just icons in a subjective interface, perhaps that is no wonder.

So how can we break through our subjective perception and find objective reality? I don't know. But my collaborators and I are currently trying to solve the hard problem of consciousness by building a theory in which the underlying reality emerges from a vast network of interacting conscious agents and their experiences. Our space-time interface – together with shapes, colours and other sensory properties – is as a visualisation tool that some agents, like us, use to simplify and interact with this network.

Our hypothesis, of course, is probably wrong. But the point of science is to be precise, so we can find out precisely what is wrong with the idea. Our theory of interacting conscious agents fails if its predictions don't square with well-tested results of classical physics, quantum theory, general relativity, evolution by natural selection and so on in our space-time interface.

And the argument turns on itself. We used the theory of evolution by natural selection to discover that what we perceive isn't objective reality, but an interface with it. Now we realise that evolution itself may be just an interface projection of deeper dynamics stemming from a network of conscious agents. The goal ahead is to work out those dynamics in detail, and figure out how, precisely, they map onto our space-time interface. This will allow us to make empirical predictions testable by experiments within our subjective reality.

Science so far has focused its search on this immediate reality. What it has found can guide our theories and test our predictions as we try to look beyond it, to find the nature of objective reality. Can we do it? Just like I take out life insurance, I'm betting we can. ▮