DON HOFFMAN INTERVIEW PART TWO

Hello again, Ars Technica listeners. This is the second installment of a three-part interview with UC Irvine quantitative psychologist Don Hoffman, and his wildly original, and quite mind-boggling take on the nature of root reality.

If you haven't yet heard part one, there's a link on the page where this player's embedded, and I strongly suggest that you go back and listen to it before this one.

And with that - back to my conversation with Don Hoffman.

TRANSITION MUSIC

Rob Reid: Let's now get into the theory that with all the hacks and horsepower in our visual cortex, and with different evolutionary pressures, different things could be constructed. Let's start by maybe talking about what you call the interface theory of perception.

Don Hoffman: Our assumption that evolution by natural selection shaped our perceptions to show us the truth, that theory I claim is wrong. What evolution has done instead is much more like shaping the interface on your computer, which as we talked about is there not to show you the diodes and resistors and voltages and magnetic fields. It's there to hide the truth and to simplify the reality in a way that you can use it to do what you need to do. It shows you a little paintbrush so you can paint. It shows you a little email icon so you can make emails. It shows you a trash can ... The relationship with the reality is not one of resemblance, it's not one of similarity. There are causal relationships of course between your interface and the objective reality, and that's why you can control the reality. But what you're seeing is utterly unlike that reality.

Don Hoffman: So the claim is that when we perceive space and time, that's just our desktop. We're not seeing a preexisting space and time. The idea that space has existed and was the preexisting stage before any life and before any consciousness, that's wrong, according to this interface theory. Space is something that you create right now. It's a data structure that you create for data compression and error correction information about fitness. And then three dimensional objects are the same thing. They are data structures that you create. I look over here and see a chair, I'm creating a data structure that's telling me certain fitness information. I look away, I no longer see the chair. As computer scientists say, I'm garbage collecting that data structure and throwing it away to save memory. Now I look back, and I see the chair, and I create that data structure again, because now I need that data structure.

Don Hoffman: The chair that I perceive is there to help me. It's not there to show me the truth, just like Michael Watson's basket of ivy is there to help him for example in cooking and so forth. It's not there to show him the truth. There is no real basket of ivy out there. That's silly. None of us would believe that there's a basket of ivy when you're tasting Angostura bitters. Take the lesson from Michael Watson and apply it to every physical object around you. They're no more resembling the nature of objective reality any more than Michael Watson's ivy resembles Angostura bitters.

Rob Reid: So there could be a colossally complex, nine-dimensional, let's go string theory for a moment and imagine the universe exists in nine dimensions. There could be this colossally complex nine-dimensional predatory force that's out there, and my great, great, great, great grandfather had some distortion in his brain that caused him to perceive it as giant cat. And it's no more a giant cat than these bitters are a bucket of ivy, but boy did he survive while his cousin who is gazing in awe at this nine-dimensional specter, got devoured. And that simplification got handed down. And then the interface theory of perception says that the UI is space time. As we perceive it, as Newton perceived it, as Einstein said actually it's more complicated than you think, he was still just perceiving the desktop of the computer although he was seeing it in a more fancy way than most of us. It really is just the user interface. Objects are icons, and what that means is when you look away from the chair, the chair ain't there anymore.

Don Hoffman: Exactly right. The chair has no position, no momentum, no physical properties when it's not observed, because the chair is not a pre-existing object in reality. It's a data structure that I create on the fly and destroy when I don't need it.

Rob Reid: And so if we're again to take the crossing the street idea, if that SUV is bearing down on you, if it's just an icon, why don't you just stay put, because an icon never really hurt anybody?

Don Hoffman: Great question. So I wouldn't step in front of the SUV for the same reason I wouldn't take my blue rectangular icon for an email I'm writing and carelessly drag it to the trash can. Not because I take the icon literally, the email is not blue and rectangular, but I do take it seriously. If I drag that icon carelessly to the trash can, I could lose all the work that I've done on that email. And if it was a letter or a book that I'm writing, then I'd be very, very careful, because I could lose a year of work dragging that icon to the trash can. And that's the key point. From an evolutionary point of view as you pointed out, the symbols that we have have been shaped by natural selection to help us stay alive. We better take them seriously. But just because we take them seriously doesn't mean that we take them literally. So I take the icon seriously, but I don't take it literally. We tend to jump from I must take it seriously, therefore I must take it literally. That's actually a logical flaw. That's a logical error.

Rob Reid: Going back to that proposition that I started with that we see reality as it is because that is very adaptive evolutionarily, that seems logical, when we get to the point of saying actually space time is a UI, and real reality is radically different from that, the question might be, "Why would we do such a thing? Why would we and out ancestors craft something that has no bearing to underlying reality." Your answer actually is evolutionary pressure, counterintuitively, that a species which sees reality in a pristinely accurate way competes with a rival species that throws away most or even all of reality for a more agile metaphor, the latter species will win every time. Is that correct?

Don Hoffman: That's right. So I'm saying something that's even more radical than just the idea that our perceptions are a simplification of reality. Most people would say, I understand that evolution is going to do things on the cheap and try to do things fast, and so we're going to get sketches of the table, not get the full detail of the table, and we'll just get sketches of things that we need to stay alive. I'm saying something more radical. I'm saying that when evolution by natural selection is shaping our perceptions, the selection forces are uniformly against anything at all like the truth. In evolutionary game theory, these functions tell you what payoffs you will get for the various actions you take within the state of the world, given which creature you are, given your state. The whole game is to get more fitness payoffs than your competition. So if you do anything besides look at the payoffs, you're going to lose. And it turns out, the fitness functions themselves are just utterly unlike objective reality.

Rob Reid: Now this is what you call fitness before truth.

Don Hoffman: That's right.

Rob Reid: And you've done a lot of mathematic modeling, which you argue supports this.

Don Hoffman: That's right. So with two of my graduate students, Justin Mark and Brian Marion, we did genetic algorithms and evolutionary game simulations. We ran hundreds of thousands, even millions of randomly chosen worlds with resources that we would throw into them, and we would put creatures in these worlds. We played god. We let some creatures see all the truth. We put others that saw nothing of the truth and were just tuned to the fitness functions, and we let them compete. And what we found uniformly was that organisms that saw the truth never out competed organisms of equal complexity that were just tuned to the fitness functions. They were in a world where they could wander around and forage for resources. It was a foraging game. But their perceptions had to evolve and their actions had to evolve. So their initial actions were really stupid, and their initial perceptions were completely crazy. But after 500 generations, the creatures that evolved were foraging optimally. So we could look and see of those creatures that survive 500 generations and bred, what kind of perceptions did they have? None of them saw the truth. All of them were tuned to the fitness functions. And in fact, I doubt in the complex world that at any stage in the genetic evolution would a true perceiving creature ever arise. There would be no selection pressures for it to appear at any point.

Rob Reid: Now one of the very weird elements of reality that you argue supports this is the famous double slit experiment. And for those listeners who know exactly what the double slit experiment is, lucky you. For those who do not, it's a notoriously difficult thing to explain without visuals. And in fact it's a notoriously difficult things to explain with visuals. Since I'm the non-scientist in the room, I'm going to attempt to explain it, because if I can understand it, most people probably can. This is an experiment it should be pointed out has been carried out probably thousands of times now, correct?

Don Hoffman: Yes.

Rob Reid: Highly, highly, consistent from experiment to experiment. This is an utterly non-controversial statement about the way reality works. Specifically, light, electrons and certain other things sometimes behave like a bunch of particles and at other times behave like waves. They have a strange dual nature, which becomes evident at a very, very tiny level, at the quantum level. So first to visualize the particular nature of light. Imagine you have one of those machines that shoots tennis balls, and they're all soaked in blue paint. And you're standing in front of a wall that has two large vertical slits in it, and behind the wall is a canvas, like an artist's canvas, a couple feet behind the wall. And you start firing these tennis balls. What'll happen is that most of the tennis balls are just gonna bounce off the wall, but a certain set of them are gonna go sailing on some angle through the left slit, and some of them are gonna go sailing on another set of angles through the right slit. After a certain amount of time, if you go in and you see your artwork that's resulted, you're going to see a cluster of blue spots where the tennis balls came through the left slit, and another cluster where they came through the right slit. And that is basically what would happen if light were behaving like a particle.

Don Hoffman: Correct.

Rob Reid: So in a sense, these tennis balls are behaving very particle-like, very much like lone photons should behave, photons being the tiniest indivisible units of light. Particles flying through two slits should behave like our tennis balls and create two clusters of dots. With photons we use photographic film or a digital sensor to record where they land rather than a canvas. Now to visualize what happens when light acts like a wave, imagine we have a huge deep pan of blue paint. It's so deep we can lower the wall with the two openings halfway into it. So now it's like we have two arched doorways on this little sea of paint instead of two slits. And the pan is filled right to the rim. One more drop of paint and it's overflowing. And the far edge of the pan is touching a fresh canvas. So we're gonna make so more art.

Rob Reid: Now on our side of the pan across from the canvas, we start dropping rocks into the paint at regular intervals to make a series of waves. Those waves radiate out in the form of expanding semicircles toward the wall with the two arched doorways. Now when a wave hits, most of it's stopped by the wall. But the parts that go through the doorways become two mini waves, which themselves start radiating out as semicircles. Now when these two new sets of waves meet, they're going to interfere with each other. Sometimes the crest of a left side wave will meet with the crest of a right side wave, and they'll join up and become a higher crest. Other times a crest will meet a trough, and they'll cancel out. At that point of the wave front, there will be no wave. Eventually these much more complicated wave come to the far end of the pan, and they splash over the side. Where two crests have teamed up, they'll make a higher and deeper mark. Where a crest and a trough meet, there's nothing, so they'll make no mark at all. And the result will be a banded pattern of paint on our canvas, light dark, light dark, light dark, or a banded pattern of light on our sensor if we're using light instead of paint. And here our light is acting like waves.

Rob Reid: Okay, now imagine you're a grad student replicating the famous double slit experiment for the umpteen thousandth time. If you just shine a steady light at the two slits, on the far side you're going to get the classic banding pattern of waves. So how do you get the two shotgun patterns? Well a logical answer might be to start firing the light just one photon at a time, because then you're only sending over solitary packets. And you'd think a solitary packet has to either go through the left slit or the right one. Once you've fired enough single photons to create a discernible pattern on the sensor, it's gonna be two shotgun clusters, right? Well, wrong. You actually get the banded pattern again. It's like the photons, despite going through one slit or the other one at a time, somehow choreographed themselves. They coordinated their landings on the sensor so that instead of producing the dual shotgun pattern of random particles, they made this very specific pattern which should only be made by waves of light. It's like you expected a mob of random people to act like individuals, but you've got a North Korean stadium putting on a synchronized show. It's like a conspiracy of photons that are ganging up to play a trick on you.

Rob Reid: So to figure out what's going on, you put a detector on each slit, which will detect whether or not a photon goes through it. This way, you'll know which side each photon goes through before making its contribution to the banding pattern. But now, the photons suddenly start acting like particles, like tennis balls, and they obediently create the two shotgun patterns rather than the banded interference pattern. It's like they know they're being watched, or measured to be more precise. And sure enough, if you turn off the detectors, the photons go back to making the banded patterns of interfering waves. Now we could go on for days about what might be causing this, but the cliff notes are that many decades after discovering this, thousands of the brightest minds on earth have no idea. Did I get it more or less right?

Don Hoffman: Yeah, that's a very good example. It's one of the most clear I've heard.

Rob Reid: So what that says, and I'm going to quote something that you've written, we don't passively observe a preexisting objective reality, but actively participate in constructing reality by our actions. And that's I guess how you would tie the double slit experiment to this whole argument that we create reality by looking at it. But you're talking about chairs not photons.

Don Hoffman: That's right. So the double slit experiment is one of the weirdest things in quantum mechanics, and it indicates that when you don't observe, you cannot say which slit the particle went through. You can't say what its position is, and for most particles you can't say what it's momentum is, what it's spin, all these properties. Unless you observe, you cannot say what those properties are.

Rob Reid: And those properties don't actually take a property until they are observed.

Don Hoffman: And that's the wild thing. And Einstein was worried about this. This is why Einstein didn't like quantum mechanics. He wanted a world in which particles really had positions and momenta even when they weren't observed, and quantum mechanics doesn't allow you to specify what those positions and momenta are unless you observe. So Einstein said, "Well, whatever quantum mechanics is, it's incomplete. Maybe it's not false, but it's not complete. There's really a position, really a momentum." That's called realism. That even if you don't measure particle as a real position, a real momentum, real definite values, that's called realism. And there's never-

Rob Reid: That reality exists whether you measure it or not.

Don Hoffman: Exactly right. The other concept is locality, that these properties like position and momentum have influences that don't propagate faster than the speed of light. And so local realism is the claim.

Rob Reid: And that is an intuition that would feel to a reasonably physics savvy person about as strong as the proposition that the earth is flat. It just seems overwhelmingly logical.

Don Hoffman: That's right. It seems like of course it's gonna be true, and it turns out it's false. It's been tested many, many times. There was a theorem by John Bell in 1963 that showed us how we could in principle test whether local realism is true, and local realism is false. It's been tested many, many times, all sorts of loopholes have been proposed and closed, and everything we close a loophole we still get the same effect. Local realism is false.

Rob Reid: And then John Wheeler, another great physicist who I think did a lot of collaboration with Niels Bohr and other people in the early 20th century on nuclear work, he came up with a thought experiment which make this weird thing slightly weirder.

Don Hoffman: That's right. So Wheeler has what's called the delayed choice experiment. It's the double slit experiment, but after you've shot that photon to the double slits, and you wait until it seems like the photons should already be past the double slits, and then you decide am I going to measure which slit it came through, or am I not going to do that and I'm just going to let an interference pattern occur. So I wait until after the photon should've already been passed to make that decision. And quantum mechanics says it shouldn't affect things. It shouldn't matter when you make that decision. So Wheeler's delayed choice experiment has been done, and quantum mechanics is right. You can make the choice after quote-unquote the particle should already gone through, past the slits, and you still get the same effect.

Rob Reid: So you're firing photons in this experiment from a significant distance. Light travels, what is it, about a foot per nanosecond or something like that?

Don Hoffman: That's right.

Rob Reid: And so in this experiment, you fire your photon a long enough distance that the scientific apparatus has enough nanoseconds to actually act. And it has gone through a slit, one of the two slits, or it has not, or it's done whatever weird quantum things it's gonna do. But the point is, after it's past the hurdle of the slits, then and only then is the decision effectively made to determine if the photon has just gone through a specific slit in the very recent past. That decision's made just as the photon's about to hit the sensor, and if we decide retroactively that we're going to pin it down to having gone through one slit or the other, the photon will then hit the sensor in a way that contributes to that shotgun pattern that suggests it was acting like a particle all along. Whereas if we don't measure it, it's going to hit the sensor in a way that contributes to a banding pattern. But because it's just about to hit the sensor when the decision's made, it either has to know what's going to be chosen in advance so that it can position itself to land at the correct spot and contribute to the correct pattern, or it has to rewrite its own history since passing through the slit in order to position itself correctly.

Rob Reid: And this incredibly creepy thing has in fact been done in the lab, and at least in theory if we were looking at some very, very distant object, like a galaxy that's ten billion light years away, had that experiment been done with those photons, they will presumably behave in the same manner, which means that on a photon that left its home star nine billion years ago, we can in effect impose nine billion years of history on that thing by observing or not observing something that it does in its last instant of travel from here to our detector.

Don Hoffman: That's right. John Wheeler actually proposed this kind of experiment, a cosmological delayed choice double slit experiment. So if you have a quasar that's ten billion light years away and you have a black hole between us and that quasar, or a big galaxy, according to Einstein's theory of general relativity, it bends space and you can get, if circumstances are right, a gravitational lens. From Earth, it could look like there are actually two quasars when in fact there's only one quasar, but it's an optical illusion created by bending of space.

Don Hoffman: So you can now ask, for each photon that comes to me, do I want to decide whether it came on the left side of the gravitational lens billions of years ago, or did it come from the right side, or do I want to just measure the interference pattern. So suppose I make the choice now, and I decide to measure which side it came on, and I find out that it went on the left side. That means I can say for the last ten billion years, that photon has been on a path that started from the quasar and went around the left side of the gravitational lens. But if instead today I had chosen to not measure that and just measure the interference pattern, then it would not be true that for the last nine billion years or ten billion years that photon had gone around the left side. So the choice I make today determines the ten billion year history of that photon.

Rob Reid: I should point out that gravitational lensing has been overwhelmingly and demonstrably proven. There's lots of images that astrophysicists have captures, and you can just google them. We see double images of galaxies. There's even some triple and quintuple images that are out there. So now, I think you've argued pretty vociferously that we are seeing a user interface that looks like space time. Space time is not objective reality. What is objective reality?

Don Hoffman: The right answer is I don't know, but as a scientist, I'm going to try to propose a theory and try to make it precise and make it testable and see where we go. So the theory that I'm proposing is that consciousness is fundamental. And by that I mean, conscious experience is like experiencing the taste of chocolate, the pain of a headache, the feel of velvet, the smell of a rose, the sound of a trumpet. All these things as conscious experiences, that these are not late comers in the universe. They are the foundational entities.

Rob Reid: Those experiences, or consciousness, a being's sense of presence and consciousness?

Don Hoffman: Those experiences themselves as part of a conscious experience are going to be fundamental, plus the agent that actually experiences them. These conscious agents have experiences, like the taste of chocolate, they have then decisions they can make about how they might want to affect the experiences of other agents. So it's all about having experiences and deciding what experiences you're going to pass and how you're going to influence other agents. So it's a big social network of conscious agents interacting, and it's essentially an infinite social network. It's infinitely complicated.

Rob Reid: And that is root reality in your current thinking.

Don Hoffman: That's right. So that's the proposal. It's all mathematically precise, so I have a definition of conscious agent, there's a whole mathematical apparatus here that I'm evaluating.

Rob Reid: That you're currently working on.

Don Hoffman: That's I'm currently working on with some colleagues. And then the idea is that each agent that's finite, when it's trying to interact with this infinite network of conscious agents, it's going to have to ignore most of that reality. It's just too complicated. So it will have to have its own interface. So it uses some of its experiences as a user interface to compress this vast social network into a space time data compression format, and physical objects are just the dumbed down symbols that it uses to represent some of its interactions with these conscious agents.

Rob Reid: Now part of the reason why you have decided to start with consciousness as the raw material of reality is what's often called the hard problem. Would you care to characterize the hard problem, and why that pointed you in this direction of starting with consciousness as opposed to starting with physical elements?

Don Hoffman: So the hard problem of consciousness is surprising. We know empirically that there are lots of correlations between brain activity and specific conscious experiences. So for example, my experience of color is correlated with activity in a part of the brain's cortex called area V4. And if I have a stroke in area V4 in the left hemisphere of my brain, I will lose all color experiences in the right part of my visual world. Very, very striking. Clean correlation.

Rob Reid: Very clear locus.

Don Hoffman: Very clear locus in the brain. And if you take a magnet and inhibit area V4 in a normal person, say in the left hemisphere, then while the person is watching, color will just drain away from their right visual world. Then you pull the magnet away, and fortunately the color comes flowing back into the right visual world. And we have dozens, perhaps hundreds of these kinds of correlations.

Rob Reid: They're called NCCs, right?

Don Hoffman: The neural correlates of consciousness, NCCs. And so this is uncontroversial. There are many, many correlations.

Rob Reid: This geographic mapping of this thing that we experience happens here in the brain, which I might naively think is pretty good proof that the brain is this physical thing that's generating this stuff.

Don Hoffman: That's causing it, that's right.

Rob Reid: But you somewhere else with this.

Don Hoffman: That's right. So the neural correlates are sometimes geographic places. Sometimes they're patterns of activities. There are different kinds of neural correlates that we might look at. So we have this data and this really clean data, and it suggests there's a very tight correlation between brain activity and conscious experience. And as you said, everybody has been assuming okay, that means that the brain is somehow causing these conscious experiences to arise, or maybe brain activity is somehow identical to these conscious experiences. The problem is, we've been trying now to develop a science of consciousness, and we actually now want mathematically precise theories.

Rob Reid: Of how from this physical substrate these conscious experiences arise.

Don Hoffman: That's right. How is it that sodium and potassium and calcium ions shuttling back and forth through holes in neural membranes, or hundreds of neural membranes, or thousands or millions, whatever it might be, how do you go from that description of the brain in a principled way to I'm now experiencing the taste of chocolate, without waving your hands, pulling a rabbit out of the hat. No magic. We want a mathematically precise theory. Why should this brain activity cause the taste of chocolate? And what would I have to do to change the brain activity to make it be the feeling of velvet. This is not philosophy anymore. This is science. We want mathematically precise rules that tie specific kinds of brain activity to specific conscious experiences in a principled way that makes new predictions.

Rob Reid: The hard problem is our failure despite an enormous amount of trying to figure out how from inert matter, consciousness can arise.

Don Hoffman: That's right. We have no scientific theories of the kind that I just described. Even worse, we have no idea how to get there. We have no plausible ideas, nothing remotely possible about how brain activity could cause conscious experiences.

Rob Reid: So, your response to that is, rather than say, well, there was a time for centuries, thousands of years, we were not getting very far with flight. But at some point, lo and behold, we got there. The common response to the situation you just described is well, we're pre-Wright Brothers here, we're doing all the block and tackle work and at some point, we're going to be a Kitty Hawk. You're saying, maybe we're going about it wrong. Instead of starting with matter and getting to consciousness, you're saying what if consciousness is route reality, and matter is actually a product of that?

Don Hoffman: Exactly.

END INTERVIEW ELEMENT OF PART TWO

TRANSITION MUSIC

So Ars Technica listeners - here we conclude the second installment of my interview with Don Hoffman and of course, Part three is coming tomorrow.

As mentioned before, if you can't wait to hear the rest of the interview, you can just head on over to my site, at after-on.com. Or, type the words After On into your favorite podcast player, and scroll through the episodes to find this one, which originally ran on April 30th. There you'll also find lots of episodes about life sciences - above all, genomics and synthetic biology. Conversations about robotics, privacy and government hacking, cryptocurrency, astrophysics, drones, and a whole lot more.

Or, you could just join me tomorrow, here on Ars.

OUTRO MUSIC