

Book Review and Highlights: Quantum Computing Since Democritus

Posted on September 1, 2014 by Scott Alexander



People sometimes confuse me with Scott Aaronson because of our similar-sounding names. I encourage this, because Scott Aaronson is awesome and it can only improve my reputation to be confused with him.

But in the end, I am not Scott Aaronson. I did not write [Quantum Computing Since Democritus](#). To be honest, I wasn't really even able to understand *Quantum Computing Since Democritus*. I knew I was in for trouble when it compared itself to *The Elegant Universe* in the foreword, since I wasn't able to get through more than a few chapters of that one. I dutifully tried to do the first couple of math problems *Democritus* set for me, and I even got a couple of them right. But eventually I realized that if I wanted to read *Democritus* the way it was supposed to be read, with full or even decent understanding, it would be a multi-year project, a page a day or worse, with my gains fading away a few days after I made them into a cloud of similar-looking formulae and three-letter abbreviations.

It left me depressed. I've [said before](#) that my lack of math talent is one of my biggest regrets in life, and here was this book that really made you understand what it must feel like to be on the cutting edge of math, proving new theorems and drawing new connections and adding to the same structure of elegant knowledge begun by Pythagoras and Archimedes and expanded by Gauss, Einstein, Turing, et cetera. All I could do was remember my own [post on burdens](#), remind myself that I was on record as saying that sometimes the IQ waterline in a certain area advances beyond your ability to contribute and that's nothing to feel guilty about.

I did finish the book. But – well, imagine a book of geography. It lists all the countries of the world and their capitals, and is meant to be so comprehensive that a reader could use it to plot the most efficient journey from Timbuktu to Kalamazoo, taking into account tolls, weather, and levels of infrastructure development along the way.

And imagine a very dumb person reading that book, unable to really absorb any of the facts, but at least understanding that the world is a place with land and ocean, and the ocean is very big and blue in color, and most of the countries and cities are on the part with the land.

That is the level at which I understood *Quantum Computing Since Democritus*. I didn't get as much as was in it, but more than nothing.

I think the biggest thing I got was – I had always thought of the physicists' God as a basically benevolent guy who fine tunes constants to create a world capable of both astounding complexity and underlying simplicity.

The vision I got from *Democritus* was of a God who was single-mindedly obsessed with enforcing a couple of rules about certain types of information you are not allowed to have *under any circumstances*. Some of these rules I'd already known about. You can't have information from outside your light cone. You can't have information about the speed and position of a particle at the same time. Others I hadn't thought about as much until reading *Democritus*. Information about when a Turing machine will halt. Information about whether certain formal systems are consistent. Precise information about the quantum state of a particle. The reason God hasn't solved world poverty yet is that He is pacing about feverishly worried that someone, somewhere, is going to be able to measure the quantum state of a particle too precisely, and dreaming up new and increasingly bizarre ways He can prevent that from happening.

Aaronson goes one level deeper than most of the other popular science writers I know and speculates on why the laws of physics are the way they are. Sometimes this is the elegance and complexity route – in his chapter on quantum physics, he argues that quantum probabilities are the squares of amplitudes because if the laws of physics were any other way – the fourth power of amplitudes, or whatever – it would fail to preserve certain useful mathematical properties. But in other cases, it's back to Obsessive God

– the laws of physics are carefully designed to preserve the rules about what information you are and aren't allowed to have.

Aaronson tries to tie his own specialty, computational complexity theory, into all of this. It's hard for me to judge how successful he is. The few times he tries to tie it into areas of philosophy I know something about – like free will – I'm not too impressed. But I could be misunderstanding him.

But once again, you get the feeling that computational complexity is about what information God will and won't let you have. It's a little less absolute – more “you can't have this information without doing the full amount of work” rather than a simple no – but it seems like the same principle. There are a bunch of situations in the book where Aaronson takes something we don't really know that much about and says it *has* to be a certain way, because if it were any other way, it could be used to solve NP problems in polynomial time, and there's no way God's going to let us do that.

Aaronson ties it all together in a very interesting way – with his story of how [Australian Actresses Are Plagiarizing My Quantum Mechanics Lectures To Sell Printers](#). He tells the story of how a printer company wanted to make a pun on “more intelligent model of printer”, so they made a commercial with intelligent models in the form of fashion models talking about quantum mechanics. And the particular quantum mechanics statement they made was a plagiarized quote from a Scott Aaronson lecture. And upon thinking about it, Aaronson decided that the quote they had chosen at random

was in fact the thesis statement that tied together everything he believed and was working on. The model had said:

But if quantum mechanics isn't physics in the usual sense — if it's not about matter, or energy, or waves, or particles — then what is it about? From my perspective, it's about information and probabilities and observables, and how they relate to each other.

That seems like as good a summary as any of *Democritus*, and a pretty good description of what I got out of it. I may not be as smart as Scott Aaronson, but on my good days I am right up there with Australian fashion models.

A list of passages I highlighted in my copy for being interesting, funny, or enlightening:

Can we prove there's no program to solve the halting problem? This is what Turing does. His key idea is not even to try to analyze the internal dynamics of such a program, supposing it existed. Instead he simply says, suppose by way of contradiction that such a program P exists. Then we can modify P to produce a new program P' that does the following. Given another program Q as its input, P' :

1. Runs forever if Q halts given its own code as input, or
2. Halts if Q runs forever given its own code as input

Now we just feed P' its own code as input. By the conditions above, P' will run forever if it halts, or halt if it runs forever. Therefore, P' – and by implication P – can't have existed in the first place.

I... I suddenly understand what the halting problem is. And there is a short proof of it that makes total sense to me. This is a completely new experience.

Oracles were apparently first studied by Turing, in his 1938 PhD thesis. Obviously anyone who could write a whole thesis about these fictitious entities would have to be an extremely pure theorist, someone who wouldn't be caught dead doing anything relevant. This was certainly true in Turing's case – indeed, he spent the years after his PhD, from 1939 to 1943, studying certain abstruse symmetry transformations in a 26 letter alphabet

ॐ_ॐ

You can look at Deep Blue, the Robbins conjecture, Google, most recently Watson – and say that's not *really* AI. That's just massive search, helped along by clever programming. Now this kind of talk drives AI researchers up a wall. They say: if you told someone in the 1960s that in 30 years we'd be able to beat the world grandmaster at chess, and asked if that would count as AI, they'd say of course it's AI. But now that we know how to do it, it's no longer AI – it's just search.

The third thing that annoys me about the Chinese Room argument is the way it gets so much mileage from a possibly misleading choice of imagery, or, one might say, by trying to sidestep the entire issue of *computational complexity* purely through clever framing. We're invited to imagine someone pushing around slips of paper with zero understanding or insight, much like the doofus freshmen who write $(a + b)^2 = a^2 + b^2$ on their math tests. But how many slips of paper are we talking about!

How big would the rule book have to be, and how quickly would you have to consult it, to carry out an intelligent Chinese conversation in anything resembling real time? If each page of the rule book corresponded to one neuron of a native speaker's brain, then probably we'd be talking about a "rule book" at least the size of the Earth, its pages searchable by a swarm of robots traveling at close to the speed of light. When you put it that way, maybe it's not so hard to imagine this enormous Chinese-speaking entity that we've brought into being might have something we'd be prepared to call understanding or insight.

This is a really clever counterargument to Chinese Room I'd never heard before. Philosophers are so good at pure qualitative distinctions that it's easy to slip the difference between "guy in a room" and "planet being processed by lightspeed robots" under the rug.

Many people's anti-robot animus is probably a combination of two ingredients – the directly experienced certainty that

they're conscious – that they perceive sounds, colors, etc – and the belief that if they were just a computation, then they could not be conscious in this way. For people who think this way, granting consciousness to a robot seems strangely equivalent to denying that one is conscious oneself.

This is actually a pretty deep way of looking at it.

My contention in this chapter is that quantum mechanics is what you would inevitably come up with if you started from probability theory, and then said, let's try to generalize it so that the numbers we used to call "probabilities" can be negative numbers. As such, the theory could have been invented by mathematicians in the nineteenth century without any input from experiment. It wasn't, but it could have been. And yet, with all the structures mathematicians studied, none of them came up with quantum mechanics until experiment forced it on them.

Aaronson's explanation of quantum mechanics is a lot like Eliezer's explanation of quantum mechanics, in that they both start by saying that the famous counterintuitiveness of the subject is partly because people choose to teach it in a backwards way in order to mirror the historical progress of understanding. I'm sure Eliezer mentioned it many times, but I didn't really get the understanding of amplitudes as potentially negative probability-type-things until I read Aaronson.

And that's a perfect illustration of why experiments are necessary in the first place! More often than not, the only reason we need experiments is that we're not smart enough. After the experiment has been done, if we've learned anything worth knowing at all, then we hope we've learned why the experiment wasn't necessary to begin with – why it wouldn't have made sense for the universe to be any other way. But we're too dumb to figure it out ourselves

Compare: [Einstein's Arrogance](#), [Negative Creativity](#).

Quantum mechanics does offer a way out [the philosophical puzzle about whether you “survive” a teleportation where a machine scans you on an atomic level, radios the data to Mars, another machine on Mars makes an atom-for-atom copy of you, and then the original is destroyed]. Suppose some of the information that made you you was actually quantum information. Then, even if you were a thoroughgoing materialist, you could still have an excellent reason not to use the teleportation machine – because, as a consequence of the No-Cloning Theorem, no such machine could possibly work as claimed

This is fighting the hypothetical a little, but maybe in a productive way.

[Bayesianism] is one way to do it, but computational learning theory tells us that it's not the only way. You don't need to

start out with an assumption about a probability distribution over the hypothesis. You can make a worst-case assumption about the hypothesis and then just say that you'd like to learn any hypothesis in the concept class, for any sample distribution, with high probability over the choice of samples. In other words, you can trade the Bayesians' probability distribution over hypotheses for a probability distribution over sample data.

I hear a bunch of people telling me Bayesianism isn't everything, it's the only thing – and another bunch of people telling me it's one useful tool in an entire bag of them. I didn't understand enough of the book's chapter on computational learning to gain too much insight here, but I will tick off one more name as being on the "one useful tool" side. Also, it makes me angry that Scott Aaronson knows so much about computational learning theory. He already knows lots of complicated stuff about computers, quantum physics, set theory, and philosophy. Part of me wants to get angry: WHY IS ONE PERSON ALLOWED TO BE SO SMART? But I guess it's more like how I know more than average about history, literature, geography, etc. I guess if you have high math ability and some intellectual curiosity, you end up able to plug it into everything pretty effortlessly. Don't care though. Still jealous.

Imagine there's a very large population of people in the world, and that there's a madman. What the madman does is, he kidnaps ten people and puts them in a room. He then throws a pair of dice. If the dice land snake-eyes (two ones) then he murders everyone in the room. If the dice do not

land snake-eyes, then he releases everyone, then kidnaps 100 new people. He now does the same thing: he rolls two dice; if they land snake-eyes, he kills everyone, and if they don't land snake-eyes, then he releases them and kidnaps 1000 people. He keeps doing this until he gets snake-eyes, at which point he's done. So now, imagine that you've been kidnapped. Conditioned on that fact, how likely is it that you're going to die? One answer is that the dice have a $1/36$ chance of landing snake eyes, so you should only be a "little bit" worried (considering). A second reflection you could make is to consider, of people who enter the room, what the fraction is of people who ever get out. About $8/9$ of the people who ever go into the room will die.

This interested me because it is equivalent to the Anthropic Doomsday conjecture and I'd never heard this phrasing of it before.

Finally, if we want to combine the anthropic computation idea with the Doomsday Argument, then there's the Adam and Eve puzzle. Suppose Adam and Eve are the first two observers, and that they'd like to solve an instance of an NP-complete problem, say, 3-SAT. To do so, they pick a random assignment, and form a very clear intention beforehand that if the assignment happens to be satisfying, they won't have any kids, whereas if the assignment is not satisfying, then they will go forth and multiply. Now let's assume SSA. Then, conditioned on having chosen an unsatisfying assignment, how likely is it that they would be an Adam and Eve in the first place, as opposed to one of the vast number of future

observers? Therefore, conditioned upon the fact that they are the first two observers, the SSA predicts that, with overwhelming probability, they will pick a satisfying assignment.

And the Lord saw Eve and said “What are you doing?”. And Eve said “I am forming an intention not to reproduce if I generate a solution to an NP complete problem, as part of an experiment in anthropic computation”. And the Lord asked “Who told you this?” And Eve said “It was the serpent who bade me compute, for he told me if I did this I would be as God, knowing subgraph isomorphism and 3SAT.” Then the Lord cast them forth from the Garden, because He was Information Theoretic God and preventing people from screwing with complexity classes is like His entire shtick.

I like to engage skeptics for several reasons. First of all, because I like arguing. Second, often I find that the best way to come up with new results is to find someone who’s saying something that seems clearly, manifestly wrong to me, and then try to think of counterarguments. Wrong claims are a fertile source of research ideas.

I said something almost exactly the same on Facebook a few days ago when Brienne asked how to generate good ideas.

There’s a joke about a planet full of people who believe in anti-induction: if the sun has risen every day in the past, then today, we should expect that it won’t. As a result, these people are all starving and living in poverty. Someone visits

the planet and tells them, “Hey, why are you still using this anti-induction philosophy? You’re living in horrible poverty!” They answer, “Well, it never worked before.”

ಠ_ಠ