Rationally Speaking #143: Scott Aaronson on, “The theorem that proves rationalists can’t disagree”

Julia: Welcome to Rationally Speaking, the podcast where we explore the borderlands between reason and nonsense. I'm your host, Julia Galef, and with us today is our guest, Professor Scott Aaronson.

Scott is a professor of Electrical Engineering and Computer Science at MIT. He blogs at *Shtetl Optimized* and he's also the author of the book *Quantum Computing Since Democritus*, which our listeners might recognize because a recent Rationally Speaking guest, physicist Sean Carroll, named it as his pick of the episode and called it, "This generation's *Godel, Escher, Bach*," which I think is the best introduction a guest could possibly get on Rationally Speaking.

Scott, welcome to the show.

Scott: Sean was too kind, but thank you so much, Julia. I'm a big fan of your work, so it's an honor to be on your podcast.

Julia: Wonderful. Let's tell our listeners what we're going to talk about today. Recently, Scott was in Berkeley and he visited the Summer Program on Applied Rationality and Cognition, or SPARC for short, which is a summer program for gifted high school math students that my organization runs. Scott gave a guest lecture, which is also posted on his blog. You can read the transcript there.

The topic of the lecture was, basically, rational disagreement – so, what should it look like when two rational people have a disagreement and share their opinions and discuss? For example, should we expect their opinions to converge by the end of the conversation? Does being rational -- and do the rules of rationality -- allow for disagreements to persist, basically?

That's my brief introduction to the topic. Scott, I'll turn the floor over to you. You can give an assuredly better introduction.

Scott: I don't know. Aumann's theorem is a theorem in interactive epistemology, which I guess is a field that didn't really exist before the theorem came along. It was proved by Robert Aumann in the 1970s. It was part of what he won the Nobel Prize in Economics for. I think, like many things in economics, once you have the right idea or you know what you're talking about, it's a few lines to prove it. Almost all the work is just adopting a certain point of view.

What this theorem says is basically that if you've got two rational agents who are Bayesian, which is that they obey, they both believe, the rules of the probability calculus; and they've got common priors, which means that they assign the same probabilities before seeing anything to each possible state of the world; and we all know that they're honest and rational; and we all know that we know that, and know that we know that we know that, and so forth. These are admittedly big assumptions.

Then these agents go off and live separate lives. They have different experiences. One of them reads one book, one of them reads another; one visits one country, and so forth. They have different experiences, which then causes them to update their beliefs differently.

Then they meet again. And they're discussing some question, which could be anything that depends on, is a function of, the state of the world. Like, what is the probability that Hillary will win the election? Or what's the probability that there is extraterrestrial life.

What it says -- of course having had different experiences, having different knowledge, they may have different beliefs about this question. But it says, if they can ever reach a position where their opinions are common knowledge to each other, meaning that they both know each other's opinions and they both know that they know them, and they know that they know that they know them and so on, forever -- then that very fact implies that their opinion *has* to be equal. It's often summarized by saying that they cannot agree to disagree. Having common knowledge of each other's opinions precludes the opinions being different.

It's a surprising theorem when you first hear it, because it seems to make a prediction about how rational agents are going to behave, which is just wildly at variance with anything that we've ever seen in our ordinary lives. And so then a large part of the discussion around this theorem is about how do you --

Julia: -- how do we reconcile that.

Scott: Yeah, exactly, exactly. There are two ways you could use this theorem. You could view it as making a prediction about what disagreements are going to look like. And in that case, it is very clearly a failed prediction. We can just throw that out right away.

Then, you could also view this theorem as having an aspirational significance, as saying, this is what disagreements *should* look like. This is what they would look like if we were honest and rational, that actually, disagreements would not persist.

I should mention, there are a bunch of results in the vicinity of Aumann's theorem related to it somehow. One of them says if you've got rational Bayesians with common priors and so forth, then they're never going to exchange in speculative trade. They're never going to buy or sell stocks. It's just out of speculation ...

Julia: Because if I want to buy your stock or if I want to sell mine, then that gives you some information -- that reveals that I know something that you don't, that if you knew it would make you not want to do the deal?

Scott: Yep! Yeah, exactly, exactly. Also, suppose we have a conversation about some question where I tell you my opinion and then that causes you to change your opinion. Just the mere fact that I'm an honest, rational person and I believe such and such, even without

knowing any of the evidence on which my opinion is based, that may cause you to have a new opinion.

Okay, and I suppose that you would tell me your new opinion. Then, you being honest and rational, of course I update on that, and so we keep going back and forth like that. Then you can prove that under very, very general conditions, that process must terminate after a finite number of steps with our having the same opinion, assuming there are the common priors and then all of that. Furthermore, it would be common knowledge that we had the same opinion.

Now, if you want a really weird aspect of this picture, it's that while we're having this conversation, I should not be able to predict the direction in which you're next going to disagree with me. For example, if I said a very liberal opinion and you responded with a more conservative opinion, and then maybe I saw some of your points so I moved a little bit in your direction -- but I should still expect that after I respond, the next thing you're going to say is going to be more on the conservative side, right?

Julia: Right, so maybe I will have moved a little bit towards your position? Maybe, although that doesn't usually happen, but you would be very surprised if I had updated all the way towards your opinion or even past it.

Scott: Exactly, exactly. We would gradually inch towards each other, the ideal case, but I can still predict that the next thing you're going to say is going to be to my right, you can predict that the next thing I say will be to your left. That is not how a Bayesian conversation with the common prior should work at all.

Julia: Why is that? Can you give the intuition?

Scott: Yeah. The intuition for that is basically that the conversation should satisfy what's called the Martingale property, which means I am trying as hard as possible to get the right answer. Once you tell me something, I'm going to just update on that immediately. It's not grudging, it's not bargaining, "I'll grant you this if you grant me that." I'm just trying to converge to the truth.

I will update. And then, if I could predict that the next thing that you were going to say would tend to move me in a conservative direction, all else being equal or on average, whatever, then I would've *already* moved in that conservative direction anyway before I even heard your message.

Julia: Since we're giving the example of using liberal and conservative political views, I want to disambiguate that the kinds of disputes that Aumann's theorem applies to -- correct me if I'm wrong, but – they're factual questions. Not, for example, values or preferences, which we do not necessarily expect people to converge on.

Scott: Right. It's a formal framework, and it applies anytime your opinions are functions of the state of the world, where the state of the world is the thing that we have a prior over.

What that means is that if we both agreed on having a Bayesian prior probability distribution over possible moral beliefs about abortion, or about whatever, and if you and I had the same prior, then we could have an Aumannian conversation about that. Just as well as we could have one about a factual question, what's the probability of there being extraterrestrials.

In practice, it might be even weirder to imagine that we would start out with prior probabilities over moral things than over empirical ones. But it applies in any situation where we're Bayesians with common priors.

Julia: It's actually, when you think about it, striking how un-striking it is that disagreements don't resolve like this, right?

Scott: Yeah.

Julia: It's very, very normal and familiar to have a disagreement between two people who are both smart and respect each other, respect each other's thinking skill and knowledge and expertise, and they also are entering their conversation in good faith, like they're actually trying to listen to each other.

And it just goes for a long time, they share their data and their opinions, and by the end of the conversation, they haven't convinced each other, or maybe they have convinced each other a little bit, but not very much. And that's just very normal. That's not surprising at all, normally, until you view it in the context of this theorem.

Scott: I think in such cases, it's like each one wears their opinion as a badge of their identity and they're agreeing to disagree. It's one person saying, "I respect that you holding this point of view is part of your self-presentation. You respect that my holding this other point of view is part of mine," but this is not the behavior of ideal truth seekers certainly.

Julia: I agree that's a very normal process to see as well, but I was actually trying to point at something even stronger. Which is, even when people aren't (at least consciously) holding their opinions as badges of identity, even when they feel like they're trying to be completely objective and evaluate all the facts and just come to the right answer together, they still usually don't converge. On large questions, at least. If it's a small question, like, “Is the office likely to be open on Sunday?”, then maybe they'll converge on that, but not bigger questions.

Scott: It is striking. I said in my blog that the closest that I've ever seen to the prediction of Aumann's theorem is probably when you've got two mathematicians or scientists who are arguing about something in front of a blackboard. I've seen it happen, or sometimes I've been one of the participants, that two people will just be vigorously arguing about something and then 10 minutes later will still be vigorously arguing -- except now the sides have switched.

Julia: That's right, that's right.

Scott: ... Because now one person says, "Fine. I get it. This conjecture is probably true," and the other one says, "No, but maybe there is a counter-example."

Julia: Can you imagine that happening with politics? That's laughable just to even imagine that situation with politics.

Scott: That rapid switching of the direction of disagreement, that's exactly what an Aumannian conversation should look like, because these are almost exceptions that prove the rule. Certainly these are cases where the two people come to the conversation with an enormous shared background of assumptions of what's going to count as an answer and where there is a right answer.

There is an underlying ground truth. And they both hopefully understand the terms being used, that they're talking about the same thing. Even there, this is very difficult, but so much the more so in most of the other things that we argue about.

Julia: To try to reconcile the predictions that Aumann's theorem makes with what we actually see among humans, even humans that are smart and arguing in good faith, et cetera, it seems like you could go at least two different directions there.

One is to say humans aren't rational, they just aren't ---

Scott: They would say, "Everyone besides me is irrational."

Julia: Oh, right, yes, that's the more precise way to phrase the complete-

Scott: Me and you.

Julia: Right. Thank you.

Scott: Unless you and I find something we disagree about, then-

Julia: In which case it's just you.

Scott: Yeah.

Julia: So that's one thing, is to say: even if people think they're trying to update rationally, et cetera, et cetera, they're not really doing that.

And another way is to say there's some condition of Aumann's theorem, part of the formal specifications under which it holds, that doesn't apply to real life for some reason.

Aumann's theorem says people will converge with a finite amount of time, but it doesn't say how long that finite amount of time is. And maybe it's just far, far longer than the hours that humans could possibly spend discussing, and that's why we don't see convergence.

Or maybe the condition that doesn't hold is about having common priors. Like, people obviously grow up in different environments, they have different information, maybe we can never fully share it.

… Actually, I'm not quite clear on what it means to have common priors in the context of these discussions.

Scott: We can talk about all these things.

When I first learned about Aumann's theorem as a grad student in 2003 or so, my first reaction was, this has to be nonsense. There has to be some loopholes that prevent this from having any significance to the real world, even sort of aspirational.

The first thing that occurred to me was, how long did this conversation take? Exactly as you said, if it were going to take 10 billion years for us to reach common knowledge, if essentially we would need to mind-meld, or just I would need to share everything I had experienced in my life, you would need to share everything you'd experienced in yours. Of course, at that point there would be nothing to distinguish us. Not surprisingly, we would have the same beliefs, but that's just not practical. We have these very, very limited communication channels: our mouths, our ears...

This was my one little contribution to this literature that was very surprising to me: That this is not the case at all. That actually, if you just want two people to approximately agree to within 10%, let's say, about the probability that Hillary will win the election, and you want that to happen with high probability over their shared prior, so you want the two people to come into the conversation expecting there's a 90% chance that they're going to approximately agree, then this only requires a very short conversation.

The number of times we have to tell each other our opinion actually only depends on how closely we want to agree and with what probability. It doesn't even depend at all on the total amount of knowledge that we have. That was my complexity version of Aumann's theorem.

Julia: Interesting.

Scott: The other way out is that you could reject this common priors assumption. Let's go into a little bit what that means.

In Bayesian probability theory, which is maybe the most famous modern conception of what it means to be rational, you just assign these numbers, probabilities to different possible ways that the world could be. These are allowed to start out as just made-up numbers in your head, as long as they obey the rules of probability. Like that the probabilities for mutually exclusive alternatives should add up to one, and so forth.

Then, Bayes rule is this very particular rule for how you should update these numbers when you get new information, but of course, it doesn't tell you how you should come up with the numbers in the first place. Before you've seen anything, what should be

your prior probability that extraterrestrials would exist amongst all the possible universes in which you could find yourself? In what fraction of them are there extraterrestrials?

That can become almost a metaphysical question, and the traditional answer is that you're allowed to start with whatever prior probabilities you want. Maybe you can critique some as being better than others, but the math will operate with any priors equally well. The hope is that once you've done enough conditioning, then even people who started out with very different prior probabilities could still come to agree with each other.

The truth is, if you drop this assumption of common priors, meaning you say, "I come into the world with different prior expectations than you come into the world with," then Aumann's theorem doesn't work anymore. We could, in that case, agree to disagree.You see, you might say then maybe that's the obvious way out of this.

But there's a paper by Tyler Cowen and Robin Hanson, which is the one that first introduced me to this stuff, actually, called "Are Disagreements Honest?" and they have a very striking answer that they suggest to this. Which is, if we're really going to be thoroughgoing Bayesians, then you could say, how does it even make sense for two people to have different priors?

What they say is, if two people are disagreeing, say, Alice and Bob, then Alice ought to treat the very fact that she is Alice and not Bob as just another event in her life, another observation for her to conditionalize on. And Bob should treat the fact that he is Bob and not Alice as just something that he has observed.

It's a very Spinozan point of view -- that you can imagine, before any of us were born, we were pure reasoning minds, behind some veil of ignorance. In this original condition, there's nothing to differentiate any of us. Anything that did differentiate us would just be information that we would condition on.

Julia: Right, and you don't actually need the metaphysical setup in order to see the force of the argument. You could just say-

Scott: That's right, that's right. I'm telling the story, but we could restate it by saying, look, why should the fact that I am me, and not you, give my prior any sort of special epistemic privilege over your prior?

Julia: In other words, why should I believe my opinion just because it's mine?

Scott: Yeah. That's right, that's right.

Julia: Basically, in this argument, people's opinions are all just sources of data about what the true answer might be, and that goes for your opinion as well as other people's opinions.

Scott: That's right.

Julia: Do you have an answer to that? It feels very intuitively hard to accept.

Scott: Yes. One can point out that this sort of vision of these pure reasoning beings that emerge from the womb and then start conditionalizing, start updating, this is not a very good picture of what humans are. Even to learn the rules of Bayes' theorem, that takes a lot of education and so forth. You can try to point to the metaphysical strangeness of it all.

My basic point of view is I am extremely happy to be a Bayesian. Like, I believe in Bayes' theorem, like I believe in the Pythagorean theorem or whatever, and I will happily use it in any situation where it seems sensible to me to discuss a prior. If we're talking about red balls and blue balls in an urn and I pick one out, it seems very reasonable to say I start with a uniform prior over all of these balls, and then if I find that one is blue and then I conditionalize on it being blue and so forth.

When you get to things like, “What is the prior probability that extraterrestrial life exists?” or “What is the prior probability that I would've been born as Scott Aaronson and not as Julia Galef?” or that I would've been born in the 20th century and not at some other time, then either situation is where my confidence wavers. I no longer feel like it necessarily makes sense to talk about a prior. I could make one up, but then there's no reason why I should trust the numbers as being meaningful.

I'm a Bayesian in the sense that I will very happily *apply* this framework. I think it's a powerful framework. Whenever you can give me an argument that, yes, there's a prior here that makes sense -- but you have to give me that argument for a prior. I don't feel like I come into the world with a pre-existing prior over every possible sequence of perceptions that I could've had or something like that.

Julia: It actually seems to me that the arguments that Tyler and Robin made about, basically, "Why believe your opinion any more than someone else's just because it happens to be in your body and not their body?" -- it doesn't seem to me that that argument really relies on Aumann's agreement theorem.

Scott: Yeah. The truth is, I think that Aumann's theorem can underscore for us the value or the importance of certain epistemic practices. We probably could've realized they’re a good idea even without Aumann's theorem, but if it helps us, then so much the better.

Yes, always try to consider the other person's point of view, always try to think very hard about, “If I had been born as this other person, if I had experienced what they had experienced, then would this make more sense to me, or would I still be making the argument that I'm making now?”

I think it's very good as an epistemic practice to only put forward arguments if you feel that they can survive that sort of scrutiny. You might be wrong about it, they might still not, but you should at least try to subject them to that kind of scrutiny.

Julia: I do feel like Aumann's theorem adds value to our understanding. For example, because the common sense advice of “consider other people's opinions” -- it doesn't have nearly as much force without that framework behind it. Of, if you're not going to update, you need some kind of justification.

If you think that this person is just as smart and reasonable and knowledgeable as you, and you're not updating towards their opinion, and you both have common knowledge of each other's opinions -- then there's something that should bother you about that, basically. I think Aumann's theorem does go beyond the common sense "You should consider other people's opinions.”

Scott: I'm going to update, Julia. I think you're right about this. I think you're right that without Aumann's theorem, you wouldn't have the clear model that is so wildly different than what you would've expected intuitively, but has a sort of aspirational value that, yeah, we might never be able to reach the point where our disagreements look like this.

If we could reach the point where I am arguing for a more liberal position and then five minutes later I switch or I'm now arguing for the more conservative position, and then I switch back and so forth, that is actually a good sign. That's a sign of people updating as rapidly as the math says that they should.

Julia: Certainly, neither I nor my friends or colleagues, even though the most talented, rational ones are not perfect Aumann updaters or Bayesian updaters, but I think that this framework has influenced, to some extent, the way that I and my friends have conversations.

For example, it will often happen that, say, I'll share an opinion and then my friend and colleague Kenzi will hear it but not be that convinced, and then I'll say, "Actually, the fact that you weren't convinced by that makes me less confident in that opinion."

We'll update on the degree to which each other has updated on our opinions. We don't have the full infinitely iterated common knowledge of my beliefs about your beliefs, about my beliefs, et cetera, because we have limited computing power and time, but there's some of that going on. And it seems helpful.

Scott: I should mention, there's a very funny irony. One of the main exponents of Aumann's theorem or the person who taught me about it was Robin Hanson, and Robin is famous -- or infamous -- for having zillions of beliefs that just completely wildly go against the consensus of all the other smart educated people. You’d think, if Robin himself were an Aumannian, then shouldn't he be more worried about that?

Julia: I never actually asked him that. I'm curious what he would say. He was also a guest on Rationally Speaking, after Sean Carroll, and made many claims that I'm sure our listeners found surprising and not intuitive at all. Like the fact that people don't buy medical care because they want to get healthier, for example.

Scott: I think the most surprising thing that Robin could do would be to say something unsurprising. It's one of those classic paradoxes.

Julia: I want to just expand briefly on something I said a few minutes ago ...

Scott: Sure.

Julia: ... which is that I thought that Tyler and Robin's argument was very counter-intuitive, this argument being "Why should I just believe my opinion just because it's mine?"

The thing that I think is counter-intuitive there -- in one sense it does make sense, but I think when you carry that logic out a few steps, you get to the point where, all these billions of people in the world believe in a God, should I update on that? Is the fact that I don't believe in a God not really worth taking more seriously just because it's my own opinion?

You can justify not updating on all these billions of theists by saying “I can see the flaws in their logic” or “I can see why they have this mistaken impression.” But then the next step after that is to say, “If I were them, if I'd been born in their situation, I would have arguments against my current position, or the real me's current position.”

Once you start going down that path, it's not really clear how you can justify not just updating all your opinions to the same muted shade of gray that's just an average of everyone in the world's opinions, you know?

Scott: Yes, yes. No, I think you bring up something hugely important, which is what's sometimes called the problem of the reference class. What is your epistemic community? What is the set of people who you are going to regard as sufficiently honest and rational, that you should worry about it if they disagree with you?

On the one hand, you could say, if almost everyone throughout history has believed in a God and you know, then that should worry you. On the other hand, if you hang out with the people you spend all your time with, I don't know, rationalists or whatever -- for me, they're computer scientists -- if almost all of them don't believe in a God, or only very nominally believe in a God, and if I consider that my community, then I should be worried if I *do* believe, right?

Julia: Right, but it's only my community because I happened to be in this environment and in this group.

Scott: Maybe it's a little different for you and for me. Because I think you joined what's specifically a community of aspiring rationalists, and I joined, I guess, a community of theoretical computer scientists. Not because of their beliefs about God or whatever. That's just a byproduct. I think there always is this question of who ...

Julia: Who should we consider our epistemic peers?

Scott: Yeah. Right. Who do we consider our epistemic peers? I think in practice, there are almost always specific people whose opinion you weigh as heavily or more heavily than those of a billion other people.

This may relate to this other blog post that I once had about the eigen-trust or eigen-democracy, where you try to figure out this notion of which people are the ones who you should trust a lot. I think that the way that we figure this out in practice is by having some sort of equilibrium, where if someone is trusted a lot by people who we've already decided to trust for other reasons, then that adds to the trustworthiness of that person. If there was someone we trusted, but they're massively distrusted by other people who we trust, then that creates a problem that has to be resolved somehow. We're constantly looking for an equilibrium in our trust network.

I think a lot of debates that are nominally about some factual question or some moral question in the external world are really, under the surface, about the sort of which epistemic community is the one that should be trusted.

Julia: Scott, we have a few minutes left in this part of the podcast. You had mentioned before the show that you had a real-life case study in your field of quantum computing and computational complexity theory, so I was hoping you could just describe that briefly, and talk about how the disagreements have played out in that space.

Scott: Sure! The disagreements with which I have maybe the most real-life experience are about quantum computing. This is a proposed technology that would use quantum mechanics to solve certain problems a lot faster than we know how to today. The laws of quantum mechanics, as we understand them now, seem very unambiguously to allow this.

There are experimental groups all over the world that are actually trying to build this. Some of them are very optimistic that they may have useful devices within a decade or two decades or something, but this field has also engendered a lot of skepticism, including, some of it by physicists, some by computer scientists. Some of them will say this is all just completely a sham, this is something that can never work.

I've been very, very interested, maybe more than most of my colleagues have been, in interacting with these people and just trying to understand, where does the skepticism come from? Because my default position would be, I hope that the skeptics are right. Because if they're right, then that means that there is something wrong with our current understanding of quantum mechanics.

If there is really a deep reason why a quantum computer can't be built -- I don't mean that just it's too hard, that the money will run out or something like that, but it's really fundamentally impossible -- then there's really something that we don't understand about quantum physics itself.

How exciting would it be to discover that? That would be something a hundred times better than a mere success in building a quantum computer, but they don't see it that

way. Somehow I argue with people who just take it as obvious that quantum computing can never work and it's not even particularly interesting. Some of them will admit that at least it's interesting to prove why they're right, even though they know that they're right without the proof.

In any case, they see this enormous burden of proof falling squarely on the people who do believe in the standard picture of quantum mechanics. I've thought a lot about what could be behind this, and the best model I can come up with is that actually these may be people who have a radically different prior.

If you look at what quantum mechanics actually says, and the part of it that is relevant to quantum computing, what it says is absolutely mind-boggling. It says that to keep track of the state of even, let's say, a thousand particles, you need two to the thousandth power complex numbers. You need this gigantic vector of amplitudes for the different possible configurations that a system could be in. The size of the vector increases exponentially with the size of the system it's representing. Basically, quantum mechanics, if you take it at face value, says the world is exponentially more difficult to stimulate on a conventional computer than you would've expected a priori.

You can imagine having a prior over possible states of, over possible ways that the world could be. A lot of people talk about priors that incorporate Occam's razor. Simpler laws of physics should have a greater prior probability than more complicated laws.

I think pretty much everyone basically agrees about that, except that they may disagree about exactly what counts as simpler. But now in addition to Occam's razor, there are people who would have a computational element to their prior. I've called this "Occam's razor with complexity aftershave," where they would say, look, if your postulated laws of physics involved a huge amount of computation, then that very fact makes them massively unlikely.

The more computation you're positing would need to be done to calculate the predictions of your theory, the less likely your theory is. By that standard, if classical computing is your definition of computation, then you're going to say that any theory that allows quantum computation is just massively disfavored right off the bat.

It doesn't matter what the physicists have been saying for the last hundred years. It doesn't matter what all the results of these experiments are. Still, some out has to be found that would prevent nature from having to do these massive amounts of computation.

The thing I would point out is that if that's what you believe, then even if a scalable quantum computer is someday built, even if we actually have quantum computers that are factoring, or apparently factoring, hundred-thousand-digit numbers and breaking all the cryptography that we use on the internet -- which is one of the things that a quantum computer would be able to do -- this complexity aftershave by the skeptic would still have to say, “In our case, it is more likely that a demon stole the numbers or

a demon prevented me from making up a challenge that was as hard as I thought it was, and this computer is actually manipulating this gigantic list of amplitudes."

If you assign a prior probability of effectively zero to quantum mechanics or to the type of picture of the world that would allow quantum computation, then you've completely insulated yourself. It doesn't even matter if quantum computers are built. You still shouldn't fully believe in them.

Julia: Interesting. Have you paraphrased this model back to the quantum computing skeptics to see if they would agree that, yes, that is the framework I'm working from within? Have you tried to achieve common knowledge to see if ...

Scott: Yes. That's a tough question, because I think there are some skeptics, like Gil Kalai is a very good example, who actually tried to be very reasonable about it. Who say, "Look, I admit there's a possibility that I could be wrong and it is very exciting to just try to find out."

Sometimes that's the type of case where in spite of Aumann's theorem I feel that I can agree to disagree. I feel like as long as someone agrees that what we should be doing is trying to find out, because it's obviously a huge thing either way, then if someone agrees about that, then that's more important to me than if they disagree on the object level question, in some sense.

Then there are skeptics who I really would describe as just assigning a zero prior probability to it being possible. Leonid Levin would be a good example. But these people tend to have so many heterodox beliefs about so many things that it's hard to know where to even start, for me anyway, in having a Aumannian conversation with them.

Julia: We're almost out of time, but I wanted to congratulate you on having put your money where your mouth is and offered a, I don't remember when this was -- recently, I think -- offered a $100,000 bet that no one will be able to prove that quantum computing is impossible.

Scott: That was a while ago. That was Gil Kalai who prodded me into that. Notice what I didn't do: I'm not going to make a bet that a quantum computer will be built in X number of years. That I'm not going to do. But the point I was trying to make there is just to say that the laws of physics that allow this is very much the conservative position. There's nothing radical about it.

Julia: I appreciate the framing there. That's very nice. As, was it Tyler Cowen or, no, maybe it was Alex -- it was some excellent blogger who wrote that "A bet is a tax on bullshit." So I appreciate you contributing to our tax system.

We are just about out of time for this section of the podcast, so I'm going to wrap up the conversation and we'll move on now to the Rationally Speaking Pick.

[musical interlude]

Julia: Welcome back! Every episode, we invite our guests to choose the Rationally Speaking Pick of the episode. It's a book or a movie or a website, or whatever tickles his or her rational fancy. Scott, what's your pick for the episode?

Scott: My pick would be pretty much any of the books by one of my favorite contemporary writers, Rebecca Newberger Goldstein.

Julia: A two-time Rationally Speaking guest, in fact.

Scott: Ah! My favorite is *The Mind-Body Problem*. That's probably her most famous book, but she has a recent one called *Plato at the Googleplex*, which I also really like.

Julia: We actually did an episode on that, but not on *The Mind-Body Problem*.

Scott: I didn't realize that! I have to listen to more of your podcasts.

Just this week, she won the National Medal of the Humanities, which I thought was really cool. It's the sort of thing that I would imagine happening in some fantasy world. It's great that it actually happened in the real world.

Julia: Man, her trophy shelf must really be getting quite full at this point.

Scott: It must be overfilling. But I was into her before she was cool.

Julia: Nice. Thank you so much for joining us on the show, Scott. We'll be linking to your blog and book, as well as to your Rationally Speaking Pick, and to your talk that inspired this whole conversation. Great to have you on the show!

Scott: Great to be here! Thanks a lot, Julia.

Julia: This concludes another episode of Rationally Speaking. Join us next time for more explorations on the borderlands between reason and nonsense.