Rationally Speaking #201: Ben Buchanan on “The Cybersecurity Dilemma”

Julia: Welcome to Rationally Speaking, the podcast where we explore the borderlands between reason and nonsense. I'm your host, Julia Galef, and I'm here today with Ben Buchanan.

Ben did his PhD at King's College in London on War Studies, and he's now a Postdoctoral Fellow at Harvard University's Kennedy School of Government, where his research focuses on cybersecurity and statecraft. He just published his first book on the topic, and he's testified to the Senate. Ben, welcome to Rationally Speaking.

Ben: Thanks for having me.

Julia: Ben, your book is on something you call the “cybersecurity dilemma.” That's the title of the book. Tell us about that -- what is the dilemma?

Ben: Well, the title, *The Cybersecurity Dilemma*, comes from this older idea we have in international relations called the security dilemma.

The security dilemma, in one form or another, has been around since the ancient Greeks. And the notion of the security dilemma is that as one nation defends itself doing things that it genuinely thinks is just for its own self-defense, it unintentionally threatens other nations. And these other nations see the action and precipitate a response, and the first nation sees that and responds in its own way. You can get an escalatory spiral sometimes even towards conflict, and it's a conflict that no one wants.

What I do is, I look at how this notion applies in cyberspace, how nations defend themselves in cyberspace and what fear that causes in other nations.

Julia: So in a traditional security dilemma, if I'm understanding it correctly, you have something like a nation building up its stockpile of missiles so that it can defend itself against potential attacks, but other nations see that and worry that, "Oh, maybe they're stockpiling missiles because they plan to attack us," and so there is this kind of feedback loop.

How does this apply in cybersecurity? Like what is a thing that one could do for defensive purposes in cybersecurity that could look like an offensive move to other nations?

Ben: One of the examples that I write about is how the United States, for genuinely defensive reasons, hacked the People's Liberation Army in China in order to better prepare its cyber defenders against what it perceived to be a Chinese threat. They hacked the PLA. They looked at how the PLA was preparing to hack other targets, including the United States, and they used that information to inform the American defenses.

The problem, of course, is that had the Chinese discovered this American intrusion, this American hacking, it's deeply unlikely they would have known that it was for genuinely defensive reasons.

Julia: What does the hack look like to the recipient? Is it just like, "We can tell someone has been in our system, but we don't know who it was and what they were doing"? Or is it more specific than that?

Ben: It's often very challenging, if you're on the receiving end of a hack, to understand what the intentions were behind it. In some sense, I think the security dilemma in general is always about the impossibility of knowing with certainty or with a strong degree of confidence what's in another actor's mind, what their intentions are.

But the mechanics of cyberspace make that particularly acute and damaging, where if you're on the receiving end of a hack, and you're just looking at the computer forensics, it can be very difficult to know what the intentions were -- and not only that, but what the intentions will be in the future, if that hack is a beachhead for something more.

Julia: How is the cybersecurity dilemma different from previous instantiations of the security dilemma with past technologies like missiles, or nuclear weapons even? Is it just the same kind of calculus, but now applied to cybersecurity, or is the calculus fundamentally changed by the new technology?

Ben: Certainly there are, we might call "conceptual constancies," old ideas of the security dilemma that are still relevant in the cybersecurity context, but one of the things that interested me was the ways in which the mechanics of cyberspace, the physics of cyberspace, make the security dilemma worse.

One of these is the very strong linkage between intelligence collection, which often is done with defensive intent, and attack. It's very difficult for one human spy in a country to launch a full-out invasion, so the linkage between the intelligence collection by one spy and attacking army is usually not too great. There are some intermediary steps that are required to go from the spy to the full-on invasion.

But in cyber operations in nation-state hacking, oftentimes, on the receiving end, the intelligence collection looks very similar to preparation for an all-out attack. And indeed, that kind of intelligence collection is required if you want to launch a high-end cyber attack. That's one of the things, I think, that's particular to cybersecurity that makes the security dilemma in this instance more acute and more dangerous.

Julia: You pointed out in your book that there's this inherent imbalance in cybersecurity that's not present, or not present in the same way, with previous geopolitics. Which is that the advantage in cybersecurity goes to offensive moves instead of defensive moves. Can you say a little bit more about that, and why does that matter?

Ben: Sure. What's really important here is not just the reality of where are the advantages between offense and defense, what we call the offense-defense balance -- and I think in cyberspace, it's a little bit too soon to say how it's going to play out in the long

run. What's important is the *perception*. What do policymakers *think* is better, offense or defense?

The reason why is, if they think offense is better, historically, in past instances of the security dilemma, we've seen that this is more likely to cause conflict, because they want to seize the initiative and claim the offensive advantage.

What's concerning to me is that that's what we see in cyberspace. Where policymakers, up to and including folks like President Obama, talk about, in almost exactly these terms, the offense-dominant nature of cyberspace, how it's hard for defenders to keep up with skilled hackers. The concern is that in a crisis or in a conflict, that would be a spur to greater action and escalation. Whereas if they thought defense was dominant, they'd be more likely to wait and to try to let the other side act first in order to claim the advantages that accrue to the defense.

Julia: So if you're a state, and you notice that your system has been hacked, as you were saying a few minutes ago, it can be quite tricky to figure out who the culprit is or what they were there for. Does that uncertainty help or hurt?

Like, one argument you could make is that the uncertainty, the difficulty of attributing a particular intrusion to one actor versus another actor, could *mitigate* the problem. In that if I'm China and I get hacked, and I don't know who hacked me, who do I retaliate against? It's not clear.

Whereas in previous security dilemmas, if a state adds forces to a particular border, even if it's for defensive reasons, it's clear whose forces they are, and that's not a mystery. I could imagine the retaliatory feedback loop being stronger in that case. But how does that play out with attribution in cyber?

Ben: That's right. This is the famous “attribution problem,” of which much is made in the cybersecurity context, and I think there's two ways we can think about this. The first is that I don't think attribution is nearly as challenging as it's often made out to be. I don't think there are that many examples of significant cyber attacks that we've seen where we don't know with pretty strong confidence, even in public, who is responsible for it.

Julia: How do we figure that out?

Ben: Well, sometimes it's a matter of the computer forensics. So, for example, when Russia hacked the DNC in the summer of 2016, the computer forensics were very clear very quickly. The cybersecurity community has a very robust research component to it. A lot of these folks are former members of the intelligence community, the United States, the United Kingdom, and they do computer forensics and incident response. And often times, they'll publish the data, so there was no doubt for anyone looking at this data, by July or so of 2016, that it was Russia that hacked the DNC. So, in some cases, the evidence can be quite clear.

What's also interesting to me is one of the ways in which nations – so, not the cybersecurity research community, but nations -- resolve the attribution question, is

they hack each other to see what the other sides are up to. Sometimes, in advance of being hacked themselves.

Now, the New York Times reported, for example, that one of the reasons the United States was so certain, so quickly, that North Korea was the one that hacked Sony in 2014 was not just because that forensic evidence was pretty clear, which it was, but that North Korea had suffered an intrusion from the United States. That American intelligence hackers were watching the North Korean networks and watching the North Korean hackers, and saw them carry out the attack on Sony.

Now, the challenge there, of course, is that, when you're hacking for attribution purposes, it animates a lot of the same risks that the security dilemma would -- that a nation suffering the intrusion doesn't think that you're hacking just for intelligence collection purposes or defensive purposes. They think you're preparing an attack. So it gets us back to square one, unfortunately, when it comes to the cybersecurity dilemma.

Julia: Is the attribution problem going to get easier or harder? I can imagine technology and innovation pushing in both directions – like, people getting more innovative at figuring out who did it, but also people getting more innovative at cloaking who did it. Which of those forces do you think is going to dominate?

Ben: Sure, the cat-and-mouse game, in so many ways, in cyber operations continues, so the cat-and-mouse game will certainly play out in attribution. I don't know how it will resolve.

If I were to make a best guess, it's that attribution would be harder for low-end activities. So if you're using off-the-shelf tools, for example, you might be better able to hide in the noise. But that for more sophisticated operations, unless you're truly prioritizing operational security, that is to say, not getting caught, and not being attributed, those might become a little bit more obvious.

Because when they're found, there's not too many actors that can conduct them, and there's enough data out there about how different actors operate, that if they're taking any shortcuts, if they're re-using any code, or re-using any techniques, they're more likely to be attributed.

The same way that police look at a string of bank robberies and look for patterns between the robberies to identify who might be carrying out the robberies. And if the same group of robbers is carrying out the string of robberies or serial killings, whatever it is. That's what defenders do, and analysts do, when they look at cyber incidents. So to avoid that, you really have to maintain a very high level of operational security.

Julia: What are some of the strategies that we – “we” meaning states around the world historically -- have used to try to mitigate the dilemma, to try to prevent these unfortunate retaliatory feedback loops? And to what extent do those measures help with the cybersecurity version of that dilemma?

Ben: Conceptually, there's at least two categories of activities we can think about, that in the past have mitigated the security dilemma.

One of these we call the offense/defense balance. I mentioned that before. If you can change the perception of who has the advantage, and make everyone think the defense has the advantage, and maybe even give the defense the genuine advantage, then that has a tendency to reduce conflict.

So one of the great examples here is: the Russian railroad tie and railroad system uses a different width that the European one. This has the effect of making it very hard to move into Russia by train, and very hard to move out. So this makes it pretty clear that the border is reasonably stable, because to overcome that border would require a lot of effort. It gives the defense the advantage on both sides. That sort of reassures everyone.

Other points, of not artificial geography, but real geography -- whether it's the placement of islands, or the locations of certain cliffs, and so forth -- can lead to more stability in the system. That's often why international borders correspond to these geographic features. Sometimes, like in the Washington Naval Treaty, in the 1920s and 30s, this was a key component of the agreement between the United States and other sea-faring nations -- to try to get everyone just enough ships to protect their borders, given the geography, but not actually attack. That's one category.

Julia: Interesting.

Ben: A second category of what we call mitigation is to shape offense/defense differentiation. To make it very clear that the technologies you're building are purely defensive, and shouldn't be seen as threatening. So, you can imagine that technologies like, walls, and mines, and fortifications, are pretty clearly defensive in nature. You can't take territory with a wall the way you can with tanks or fast-attack jets.

On balance, if states prioritize building defensive weapons over offensive weapons or dual-use weapons, that has the effect of mitigating the security dilemma, reducing the tension that everyone feels.

Julia: Great. So what are the cybersecurity versions of those two strategies?

Ben: The challenge is that those two categories of mitigations, which have worked pretty well for a long time, don't translate very well to cybersecurity. As I said, the perception, already, on offense/defense is that offense has the advantage, and that could be destabilizing in a crisis.

In terms of offense/defense differentiation, as I mentioned, it's very hard if you're on the receiving end of an intrusion, to know if it's an offensive intrusion setting up for an attack; if it's something in the middle, setting up for an attack, but just deterance, not one they actually plan on using; if it's some kind of intelligence collection that maybe you're not happy about, but isn't gonna be an attack; or if it's

something that's genuinely defensive, so still intrusive, still not something you're happy to see, but genuinely done with defensive intent, and nothing more. Very hard at the receiving end to determine those. And therefore, we say that the offense/defense differentiation is cybersecurity is pretty hard to do.

Julia: But, surely there are things that the United States, for example, could do that would be purely defensive? Kind of analogous to building walls, or mines, or something like that? Like, we could train our officials in better security hygiene. We could build better firewalls, I guess. I don't know a ton about security technology, but something in the firewall space.

I'm sure that if we had unlimited money to spend on security, then it would be best to spend it all on everything, just all the defensive measures, all the offensive measures – or, not offensive, but all the ambiguous measures -- in addition to the purely defensive ones. But if the benefit of these measures is kind of fungible, then couldn't we, just in theory, decide to spend twice as much as we were spending on cybersecurity, in order to purchase the same total amount of risk reduction on our end, but of all the safe non-ambiguous kind, not the kind that might trigger a backlash if discovered?

Ben: Sure.

Julia: Does that make sense?

Ben: It definitely is the case that what I call baseline defenses, which are non-intrusive defenses, the firewalls you're talking about…

Julia: Baseline, great.

Ben: … The training, yeah, those baseline defenses can and should be purchased. For individuals like you and me, and for organizations, those are the only defensive tools available to us. It's not legal for us to hack someone else, even if we think someone else is gonna hack us soon. We've gotta make due with our own perimeter-based defenses and network-based defenses.

The real question is, in the sometimes dog-eat-dog world of geopolitics, are those baseline defenses sufficient, and…

Julia: So, basically what you're saying -- in my framework, you're saying: there might be a limit to how much risk reduction we can purchase just via baseline defenses?

Ben: Yeah, I think that's certainly what you'd hear if you brought this up to a US policymaker.

And we should be clear that the United States does hack other nations, for not-defensive purposes. It likes the capacity to have offensive options if it wants them. It likes doing intelligence collection that is intrusive. And many other nations do the same.

So all I suggest in the cybersecurity dilemma is, even if nations didn't have this intrinsic need to compete with one another, there are structural challenges that come from the physics of cyberspace that make it very difficult to get to stable outcomes. I think that's quite discouraging. It's something that in the short run, and also in the long run deserves attention.

Julia: Right. So, yeah, the dilemma is kind of premised on the idea that ... In the pure form of the situation, the states ideally would want to cooperate with each other. Like, they would want to abstain from offensive cyber attacks if they could be confident that other states were also abstaining. And then the problem, the dilemma, is just that we can't be confident of that fact.

But, it could also be the case that states have no desire to cooperate with each other. And that even if they could be confident that other states were being good and abstaining, they would still be like: "Cool, let's go do some cyber attacks ourselves."

If that's the world we lived in, then would the dilemma be all that relevant, if states were gonna do the same thing regardless of their impressions of what other states were doing? And how certain are you that we're not just in that world?

Ben: You're touching on something which is a long-held debate, almost a fundamental debate in international relations, which is: is every nation greedy? Is every state greedy?

Julia: Greedy, yeah.

Ben: Do they always want to try to gain more territory, get more for themselves? This is a debate that's gone on for decades, and I don't have any particular answer to it.

I think if you believe or if anyone believes that states are intrinsically greedy, then they tend to worry less about any kind of security dilemma, because it doesn't matter as much, when nations are always trying to seize more from one another.

If you don't believe in the intrinsic greed of nations, or at least not all nations, then that's when the security dilemma comes more into play, and that's always been the case. So, it's a branch of international relations that goes by the name "defensive realism", which focuses on things like the security dilemma. As opposed to, and I'm simplifying a little bit here, the more greed-focused side, which is known as "offensive realism".

Julia: Got it. So your position would basically be, "Look, the situation is probably somewhere in the middle, probably nations aren't purely greedy, and so the cybersecurity dilemma seems like it applies to at least some extent. We just don't know exactly how much."

Ben: Exactly, I would be somewhere in the middle.

Julia: Right, yeah.

Ben: I think it has applicability now already, but I also think that in the future -- even if we get to a point where we can show nations that greed does not pay or doesn't make sense -- we have the structural challenge that, even if we could remove greed from the equation, we've gotta find some way to bring about stability in line with what defensive realists would want.

Julia: Right, to deal with the cooperation problem, essentially.

So, in cooperation problems like this -- where everyone, in theory, wants to cooperate, but only if other people are cooperating -- one of the most important things you can do is to find ways to credibly, reliably, signal to the other players that you intend to cooperate. Are there ways for states to do that in the cybersecurity dilemma?

Ben: At best, those forms of signaling are nacent and early stage. So, we've tried to work on ways that enable communication in a crisis, to try to ratchet down the tension -- think about how the United States and the Soviet Union, after the Cuban missile crisis, adopted the idea for the red telephone. So there could be communication flows in a crisis. The United States and China, and the United States and Russia, have tried similar mechanisms in cybersecurity.

The challenge is that nuclear crises are what we call "strategic". They immediately go to the president or the senior leaders in a country. But very often, cybersecurity crises or cybersecurity operations don't make it to that level. They're what we call "operational" -- they sometimes have strategic effects, but they're at a lower level of abstraction. I think we haven't yet adapted some of our signaling mechanisms to account for that operational reality.

But, certainly, there's a case for bilateral cooperation, bilateral agreements, that do try to build credibility between nations, show them that they can trust one another, and eventually, move towards a world in which cooperation is seen as more tenable than maybe it is right now.

Julia: One thing that I was thinking of when I asked that question was: there are various things, some of which you mentioned in your book, that countries can do that would show that they care about security and defense. That are costly signals. Where states could have not done this thing, and would've gained an advantage for themselves -- and so, they're willingly sacrificing an advantage for the sake of showing that they care about cooperation, basically.

One example I was thinking of was announcing when zero-day exploits are found. Is that relevant to the cooperation problem, and if so, could you explain what it is?

Ben: Sure, it certainly is relevant. A zero-day exploit, simplifying a bit here, is a software exploit that no one else knows about, or the vendor doesn't know about. Defenses are much less likely to catch a hacker using a zero-day exploit than they are to catch a hacker using a regular exploit that's been known.

So, zero-day exploits are fairly rare and quite valuable for hackers. And the United States and other nations have very ambitious programs to both find zero-day exploits and to buy them from vendors to enable intelligence operations.

One of the things that I posit a nation could do is, when it finds some zero-day exploits, it could burn them – so, expose them to public view, so that defenses can fix them. Defenders will be able, more swiftly, to block hackers that are using them.

And this would potentially have some possibility of credible signaling. Because the United States would be giving up the capacity to launch intelligence operations with these exploits and, with said, be shifting the balance to the defense.

The question, as ever, is how much of this do you need to do in order to gain that credibility? And that's a matter of enormous debate right now.

Julia: Right. I also feel like the US has sacrificed some credibility by being... I was gonna say "deceptive about how honest it was being", but maybe that cancels out. But by purporting to be working for a collective security when actually it's only working for its own security. And when that's exposed, that's a huge hit to our ability to credibly claim things in the future.

Ben: Well, United States' policymakers would certainly tell you that the United States seeks to gather the kind of intelligence that all nations try to gather, and that the US just might be better at it than some other nations.

But you can certainly point to cases where the United States sacrificed long-run credibility for potentially short-run gains. One example is the US government has run a program for a very long time to ensure that encryption -- which protects, we use it every day to protect the communications we send online -- is secure. So, the US has a program to make sure encryption is secure.

And it verified as secure an encryption implementation that was not in fact secure, but had a backdoor in it, that, it seems, would enable the NSA to decrypt it. When this came out, obviously, it was a tremendous hit to US credibility, though, in the intranet, probably enabled intelligence operations.

Julia: Yeah.

Ben: So, there certainly are trade offs. And some would argue that that was an example of a long-run hit to credibility, that’s gonna be very hard to undo, for short-run intelligence gains. And my guess is that operation, if it was done as I described, was done because they never thought it would come out. They never thought the backdoor would be found.

Julia: Yeah, there's some line from some TV show I watched recently where a character said -- he'd just been caught doing something duplicitous, and his response was -- "Well, in my defense, I never thought you would find out."

Ben: Quite frankly, that's often how the world of international politics works.

Julia: Yeah, so, A, do you have any other ideas for how to tackle the cybersecurity dilemma and/or, B, what do you think is the most promising of the ideas we've talked about?

Ben: I think it's very challenging to make much progress on cybersecurity issues without thinking quite seriously about what the long-run state is a nation is trying to build towards.

There certainly are folks in the United States who would want to establish some series of norms, or some code of behavior to regulate how nations conduct themselves in cyber space. Saying things like, "You can't interfere in other nations' elections", or, "You can't steal intellectual property for the benefit of your own corporations."

The challenge is that very rarely is the United States or Western allies willing to admit what kind of operations they're willing to give up in order to get China and Russia to give up certain kinds of operations that are valuable to them. And I don't think it's realistic in the international system, such as it is, to expect that nations like China and Russia are gonna bend themselves to the American will without something in return.

So, I think that the biggest priority in the broader cybersecurity agenda -- beyond important things like improving defenses, which are absolutely vital -- the biggest strategic priority is figuring out what's the end state we're trying to get towards and what hard choices are we willing to make in order to get towards that end state.

Are we willing to constrain some types of intelligence collection? Are we willing to constrain some kinds of offensive preparation? Are we willing to maybe give some ground on some types of, what we call, "internet sovereignty" or "internet freedom", saying that the Chinese and the Russians can have more control over their own internets?

These are all very challenging questions, and it involves giving up things that are important to the United States and important to human rights more broadly. But I don't think there's an easy way out of the cybersecurity dilemma.

Julia: On that note, how seriously is the US government taking cybersecurity? On the one hand, my experience is that the government tends to take national security concerns very seriously. But, on the other hand, there are all these cases of departments, US government departments, getting hacked stupidly easily, where precautions were not taken that they knew should've been taken.

It's just a little hard to look at all those cases and think that the government is taking these threats seriously. You had a story in your book about the Office of Personnel Management getting hacked in 2015 that was… disturbing, to say the least.

Ben: That's right. There certainly are many cases that should prompt worry amongst taxpayers and amongst citizens, that a lot that should be done is not getting done.

That said, that type of security threat in general is definitely recognized for the severity it poses. I think every year since 2013, so going on four or five years now, the director of National Intelligence has looked at cybersecurity threats as the number one threat to United States' interests.

The way I often phrase it is the United States has the nicest rocks when it comes to cyber offense, but we still live in an exceptionally glassy house when it comes to cyber defense.

Julia: Where the glassiness of our house is… the government not taking precautions that need to be taken?

Ben: Maybe I should say "the glassiest mansion," because it's just a really big enterprise.

Julia: The place people wanna throw rocks at?

Ben: Yeah, a place people wanna throw rocks at, because there's good stuff inside -- and a place where there's just so many windows that we don't have enough defenders to patch them all up.

Julia: Yeah.

Ben: So, you mentioned the Office of Personnel Management, which is a fairly small office that happens to hold security, extend its information, for millions of Americans, and the records for tens of millions of Americans who are government employees.

This was breached, apparently, by China. And that's a real problem, but I don't know that this would have been thought of as something that needed tons of cybersecurity defense. I don't think OPM hired their first cybersecurity employee until something like 2013.

This is just one example of many of some glassy windows that weren't protected, to continue our metaphor. When you've got a big mansion, there's a lot of those.

Julia: Yes. The thing that I remember being especially disturbing about that story was that, A: some experts had been urging the department to take these measures for a long time, and nothing had happened… and B: that there had been similar hacks in the past and no precautions were taken as a result of those, and C: that there had been a previous hack attempt that had failed only because the systems they were running at the OPM were so outdated, that the hackers were stymied and didn't know how to deal with these super outdated systems. And that's why they failed to hack the department.

Ben: Yes. You know, when you're writing a book like this, you always try to figure out how can you explain things in as accessible a way as possible, and you look for telling details. I think that last one was a pretty telling detail about the state of OPM cybersecurity.

Julia: There was this Onion article a few years back with the headline, “Smart Qualified People Behind The Scenes Keeping America Safe: ‘We don't exist.’”

I think about that headline a lot. How much has your view of general institutional, or governmental, competence evolved as you've studied cybersecurity?

Ben: It's important to recognize there really are a lot of people who show up to work in the US government, for less money than they make in the private sector, on cybersecurity and elsewhere, and try to work on very hard missions.

Julia: Sure, sure, sure.

Ben: I think the challenge is we're not necessarily hiring as many of those as we need, and we have a very hard time retaining them. Things like government shutdowns make it challenging, and the general problems in terms of salary and compensation.

And so it's not that these smart competent people don't exist. It's that there's not enough of them, and frequently, they're not empowered to do what they need to do in order to achieve the mission. The mission we ask them to achieve is a very challenging one. We're talking about securing that glassy mansion. That's a very challenging mission. That probably requires many more people who are more empowered than what we've got right now.

Julia: That's a very good and very diplomatic answer, and it sounds right to me, but really, what I was wondering is: if you had to estimate what is the probability…This is a vague question that would have to be made more concrete, but something like: If you were to estimate the probability that the US government is going to "do the sensible thing" in a crisis -- or do the sensible thing to prepare for crisis, or something like that -- has your estimate of that probability gone up or down, as you've learned more about the government? From whatever your baseline was.

Ben: I think what I've come to appreciate the more I look into any kind of government apparatus is the old Washington adage, "Where you stand depends on where you sit," and that oftentimes, the missions folks are asked to achieve are parochial to their agency or to their unit. What I don't see enough of in cybersecurity is broad vision for how these different missions are going to fit together. Certainly, I think it's difficult to criticize the NSA for collecting foreign intelligence. That is their mission, but there should be some adjudication at the high level, the highest level of government.

What kind of foreign intelligence collection carry too much a risk of blow-back that we shouldn't do it? I don't think we have yet come upon a broader national cybersecurity strategy that has put the pieces together in a coherent fashion. It's not for a lack of trying. Every president in the last couple has said they wanted to make cybersecurity a priority.

There had been a number of folks in those senior jobs. Some of the work those folks has done has been very good. But cybersecurity is exciting to study because it's so

cross-cutting, but it's very challenging for governments to make policy, to make strategy in this area *because* it is so cross-cutting.

Julia: I have a question about methodology in your field. To what extent does game theory – like, the formal field of game theory -- help you analyze situations like this, like the cybersecurity dilemma?

I ask because it definitely seems like the kind of field that game theory is designed to help with. And I've always found game theory intriguing. But I worry that it may not be that useful in the real world. Because A: real world situations that involve game theory are either simple enough that a smart person with no formal theory training could probably just figure out on their own, like, "Oh, this is a cooperation problem," or B: the game theory models that do yield interesting, non-obvious answers are so stylized and over-simplified that the conditions that would make that model relevant would never actually obtain, to any significant degree, in the messy real world.

So yeah, I'm wondering how much formal game theory helps with issues like the ones you study.

Ben: I think in broader American political science, game theory certainly has an enormous purchase these days. But I tend to share some of your concerns about it. Particularly as applied to things like cybersecurity, where uncertainty is very high, where misinterpretation is very likely, that I think there are limits to how useful those models are.

I'm not a game theoretician myself, and I don't use a lot of that methodology, but I could certainly see if someone were interested in game theory, were interested in how game theory works with the security dilemma, that it would be interesting to see what they came up with when they look at something like the cybersecurity dilemma as I've read it up. Their conclusions may be similar, or they may use a model, a game theoretic model and come to a different conclusion than I've come to, and so I'd be interested in seeing if game theory leads to a different outcome or a different predictive outcome. In the long run, who gets closer to the mark?

Julia: Do you think that your approach to thinking about cybersecurity differs from other academics? I have the impression that the way that you choose topics to research, and to focus on, uses maybe some different search criteria from most academics. That maybe, for example, you're more focused on real world importance or something. I'm not sure, what's your impression?

Ben: It's always hard to compare oneself to others, but I think I'm comfortable saying that I am less quantitative and less game theoretic than the median in American political science.

But this is a Rationally Speaking podcast, so I would say that, in something like cybersecurity, you can be qualitative and not game theoretic, and still be very rational. I try to work very hard to make sure that I am drawing on technical sources that tell me, and tell everyone, what actually is happening in cyber space. I draw on

many computer forensic reports by incident responders, by cybersecurity researchers, that are not meant for political scientists. They're meant for technical researchers that are trying to improve cybersecurity defenses, but do provide an enormous window into what happens between nations in cyber space, how and why they hack one another.

I do think that American political science, American international relations scholarship, to the extent it wants to write about cybersecurity, should take those reports very seriously. And try to get over some of the technical barriers to entry, because they provide an unprecedented window into what's actually going on out there.

If our theories and our models and our quantitative methodologies don't match up with what actually is going on out there, I don't think they're terribly useful. I try very hard to start with what's happening there. And hopefully, the more formal model folks can build on this research that I've done, and apply their models and see where they come out.

Julia: A couple times throughout this conversation, you've touched on a disagreement or debate within the field. For example, there's the debate over how greedy are nations, how greedy should we model them being? Then, I guess there is also a debate about the attribution problem -- how easy is it to attribute an intrusion to one actor versus another?

Are there any other controversies or disagreements within the community of experts who study these things?

Ben: One of the big ones that I think is still evolving is how the notion of deterrence applies to cybersecurity. Deterrence in some way is what helps make international relations a field. Because in the advent of nuclear weapons, the weapons were so terrible and so powerful that the key question of policy and also of scholarship became, "How can we have these weapons but design a system that we'll never have to use these weapons?" That was the era of deterrence.

Julia: Right.

Ben: I do think that folks like Thomas Schelling and other academics in the '50s and the '60s and the '70s did enormously impactful work with the dawn of nuclear weapons, in getting humanity, getting the United States National Security Establishment in particular, to a strategy that was workable with deterrence.

The challenge, I think, is that I'm not sure the nuclear deterrence strategy applies terribly well to cyber deterrence. We often hear the policy question, "How do we deter China and Russia or other adversaries from using their hacking abilities against the United States in a way we don't like?" Getting back to what I said before about how cyber capabilities are often not strategic capabilities the way that nuclear weapons are… It's a challenge, I think, to model cyber deterrence in a new way without hewing too carefully to this well-established notion of nuclear deterrence.

So one of the debates that we're having as a field, and that i actually participate in is what other kinds of deterrence concepts could be useful when we talk about this new area of cybersecurity?

Julia: Great. Well, Ben, before we wrap up, I want to give you an opportunity to introduce the Rationally Speaking pick of the episode, which is a book or article or blog or something else you've consumed in your career that has influenced the way you think. What would your pick be for this episode?

Ben: A book that certainly I've thought about a lot is a book called *Rise Of The Machines*, which is by someone named Thomas Rid, who is a big influence in my thinking in cybersecurity. What Thomas does in *Rise Of The Machines* that's quite interesting is he focuses on the history of cybersecurity.

We think of cybersecurity as a field that, in many ways, is new and different. In some respects, it truly is, but it also has an enormously rich history going back to the 1930s and basically every decade from then until now. It is always a great reminder to read something like that, that exposes and ties together the history, because we can learn a lot from that history. It's a reminder that as we work on the problems of today, we shouldn't discard really valuable ideas from the past. Certainly, I think this security dilemma has been called an old and brilliant idea for new and dangerous times. I think the more we can focus on history, the better off we'll be to solving the cybersecurity challenges of the present and the future.

Julia: Great. Well, Ben, thank you so much for joining us on the show today. We'll link to *Rise Of The Machines* and also to your new book, *The Cybersecurity Dilemma,* on the podcast website.

Ben: My pleasure. Thanks for having me.

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