
# Applying Scientific Methods in Cybersecurity

*Featuring Leigh Metcalf and Jonathan Spring*

---

*Welcome to the SEI Podcast Series, a production of the Carnegie Mellon University Software Engineering Institute. The SEI is a federally funded research and development center sponsored by the U.S. Department of Defense. A transcript of today's podcast is posted on the SEI website at sei.cmu.edu/podcasts.*

**Suzanne Miller:** Welcome to the SEI Podcast Series. My name is Suzanne Miller. I am a principal researcher in the SEI Software Solutions Division. Today, I am very pleased to welcome Dr. Jonathan Spring, a senior vulnerability researcher in our CERT Division, and Dr. Leigh Metcalf, a senior network security research analyst also with CERT. And today, we're here to talk about applying scientific practice in cybersecurity research, which is actually the topic of a book that Dr. Spring and Dr. Metcalf have recently published, called *Using Science in Cybersecurity*. So, I want to welcome both of you today.

**Jonathan Spring:** Thanks, Suze. Happy to be here.

**Suzanne:** And we are remote again. So, we're not back in the studio yet. But I do want to start off by asking each of you to tell us a little bit about yourselves and how did you end up in this kind of work. And Jonathan, in particular, how did you end up at the SEI doing this kind of work? Why don't we go ahead and start with Leigh?

**Leigh Metcalf:** I have actually a very strange background, because I have a PhD in theoretical mathematics, specializing in a field called algebraic topology. But I also spent over 10 years in industry, working at various startups, doing various Internet-related things, and then fell into cybersecurity, mainly because no one else would do it. And I ended up at CERT, because after doing that for a while, I decided that cybersecurity was what I really wanted to do. I have been here for 11 years now, and it is a fascinating field, with lots of different problems, and I enjoy it very much. I started a journal, because I thought that researchers and practitioners weren't talking enough, and that is what we do at the Software Engineering Institute, is we sit at the intersection between research and practice.

**Suzanne:** Thank you, Leigh. And Jono, as you like to be known, what brought you here, and what is it that you like doing here?

**Jonathan:** Oh, yes. That is an interesting set of questions. I grew up in Pittsburgh. I actually was interested more in philosophy and philosophy of science when I was in university. So, I did a bunch of like biology, chemistry, physics, linguistics, that sort of stuff. But, I was graduating university in the middle of that housing crash that happened, and there weren't a whole lot of jobs for philosophers. I had also done some computer science, and so I went into a program at Pitt for information security. Some of my professors were CERT staff that were adjuncting there, and so I ended up moving another two blocks down the road once I needed a job, and going to the SEI. For a long time, I had been trying to combine the cybersecurity and philosophy-of-science stuff. I went and did a PhD, but I had to go do it in London, because even though CMU has a very strong computer-science program, and Pitt has a very strong philosophy-of-science program, those two things don't intersect a whole lot. So, with David Pym and Phyllis Illari, I was able to find at UCL [University College London] two people who were both willing to do some of those things. So, I don't know how many people have done PhDs in philosophy of cybersecurity, but…

**Suzanne:** You are the only one I have ever met. I will put it that way.

**Jonathan:** The reason for all of that is because the whole purpose of cybersecurity is to inform evidence-based policymaking. Whether it is organizational policy of what the security policy should be, or whether it is public policy, or whatever. I think that in order to do evidence-based policymaking, you have to understand how the evidence is gathered. And, if you want to understand how evidence is gathered, you have to understand how to do science, basically.

**Suzanne:** OK. That is fair. That is actually what we are here to talk about. Let's switch to that, to talking about *why*. You just really introduced us to that: how applying these methods to cybersecurity research makes it better, faster, cheaper, for practitioners, users, the government: all the many stakeholders that have an equity in cybersecurity. And, Jonathan, I am guessing you may want to take that question.

**Jonathan:** Leigh has a lot of feelings as well. We have been working together for 10 years. There is a really long history of scientific methods being just the more reliable way to gather good evidence about how the world is working. So, if you would like to know the best guess on what is going to happen if you make a change, and you are going to engineer a system or something, we want to understand what we know about it and in a reliable way. Dr. Pym has this nice piece on if we take the *cyber* part of cybersecurity seriously, what does it mean? That was coined in sci-fi books in the '80s, Gibson and *Neuromancer*, as like the overlay that we make of the social space on top of the machines. And so, if we take that seriously and we want to do

cybersecurity, we need to secure the human social spaces that we make on top of the machines as well as the machines. It makes it really clear that it is this big interdisciplinary thing that is going to require a lot of communication between technical sciences, engineering sciences, physical sciences, and social sciences. So, I think that that is its own discipline.

**Suzanne:** Fair enough. Leigh, did you want to add anything to that?

**Leigh:** Well, I think doing cybersecurity well is an important basic part of doing science and cybersecurity. And doing it well and having it repeatable actually saves time and money. So if people do it well to begin with, then moving forward, you are not repeating other people's work. You are learning from the past appropriately, because it has been done in an appropriate manner.

**Suzanne:** That is one of the things I think that the scientific methods in general have that focus on repeatability and gathering the data and evidence that allow us to repeat an experiment, to reverify a hypothesis. So, those things really come together in this area in terms of what you are saying about wanting to be able to have things that are going to save us money down the road because we know how they work. Fair enough?

**Jonathan:** Leigh that is one of the purposes behind DTRAP [*Digital Threats: Research and Practice*], the journal that you are editor-in-chief of?

**Leigh:** Yes it is. The goal was to ask the authors to, for example, tell us what data source they used for their experiment. Not just say, *I used malware*, but this is what kind of and where the malware...and where I got it from. So that someone else can come along and say, *Oh, that's how I do that. That's something I can repeat, or that's something I can use*. Methods sitting by themselves in this field, without background of what kind of data and where the data originated from aren't as useful.

**Suzanne:** Again, connecting the science, that is one of the things about science in general, is that we need to understand that the data is relevant to the problem, that it is relevant to the solution, and that it is the data that actually informs new choices that we may want to make. So that whole connection to data is a really important one in cybersecurity, as well as other elements of science. So, what is the current state of the practice on applying scientific methods to cybersecurity? We have got the magazine that gives us, what are some of the trends? What is it that you are gathering, through that source and other sources, that help us to understand what is the current state of the practice in this arena?

**Jonathan:** I am also the PC co-chair of the New Security Paradigm Workshop, this year and last year. So that I think is a smaller venue, but, you know, some more stuff like that. Let me see, it depends really on if you mean what people are calling science in cybersecurity, or what people are actually doing. The National Academies of Science, Medicine, and Engineering [National

Academies of Science, Engineering, and Medicine]—I might have got those mixed up the wrong way—about every five years, they have done a big report on what should we do for science in cybersecurity, like what should we be funding, like what should we be prioritizing. There is certainly very high-level acknowledgment that this is important, but in my reading of everything that they have written down, like the NSF [National Science Foundation], the National Academies, the White House, GCHQ [Government Communications Headquarters] over in England, the Canadians, the NSA [National Security Agency], like all of that stuff. Because the NSA has the Science of Security Symposium, all of this. They are still all pretty focused on a like 1950s conception of the scientific method, and I say the scientific method because that's what is implied there. Not like there are multiple methods for different disciplines that could all be used at an appropriate time, in different interoperating ways. But like there was one right answer, we need to try to falsify things, and like falsification, trying to prove something wrong is fine, but it doesn't tell you how to come up with a good hypothesis. It just tells you how to test one, once you have it.

**Suzanne:** Yes.

**Jonathan:** And so, a lot of what I have been interested in is describing scientific methods in a way that helps people come up with good hypotheses, a more structured way of coming up with hypotheses and how you're going to pick what thing is the next most important thing to test, rather than just saying, *Well, if it's not falsifiable, it's no good, give me another rock that's a different color.*

**Suzanne:** Well I know as a researcher myself, that idea of coming up with a good hypothesis and understanding…I like both points that you made: coming up with a good hypothesis and what is the next most important thing to test, as opposed to what's the easiest thing to test. Because I know that researchers do fall into the trap, partly to get funding. *Well, here is the thing that I can falsify. Here is the thing that I can test with a falsification hypothesis. So therefore, that is the one I am going to do.* The bigger question may be something that is not as amenable to that kind of approach. We sometimes lose that because it is perceived as not being falsifiable, and therefore not being testable in that way. So, how do you overcome that? How do you deal with that in the cybersecurity and science intersection?

**Jonathan:** Well, some of it is social, among the researchers. And so, I think that is where DTRAP is sort of well-placed. Other things like NSPW have been doing that on sort of maybe a different scale or different…

**Suzanne:** NSPW is?

**Jonathan:** The New Security Paradigms Workshop. *Paradigms*, there, very explicitly a Kuhn reference…Thomas Kuhn is the philosopher of science, after Popper, who sort of says, *No, the falsification thing doesn't make sense. Researchers work within a paradigm, where a paradigmatic example is like the example that we are working off of, and then we refine that and refine that and refine that. And then it breaks, and we've had a new one*. That is not going to happen either, but that was the...that's where NSPW is coming from. But how do we get people to use appropriate scientific methods? I mean, you have to make it worth their effort.

**Suzanne:** Sure. We have to give them answers that they can't get other ways, is one of the ways I would look at it.

**Jonathan:** And you have to make it clear somehow, when people have made a shortcut that is actually harmful.

**Suzanne:** Ah. OK.

**Jonathan:** Because some shortcuts maybe just save you time, and the loss in precision or resolution maybe is not harmful.

**Suzanne:** So, if somebody wants to bring that perspective into their work, and they are not currently using scientific methods explicitly in their cybersecurity research, what are some steps that you would suggest that they take to bring that into their work considerations? Like what are the shortcuts that are harmful versus not? How do you design good experiments, case studies, research strategies, etc., like that? What are some of the things that you would suggest to people that agree with your perspective on, we have to go beyond falsification?

**Leigh:** Well, OK. One of the major common pitfalls we see today in science is called data dredging. It is when people have a dataset, and they can't prove what they are looking for, so they keep going back and refining their guess, basically. What they end up doing is they are designing their experiment so that it fits the weirdnesses in their dataset. And data is weird. Every dataset is weird in its own way, but if you design your research to the weirdnesses in your dataset, then no one else can actually do the same thing over again. And data dredging is thought to be the reason a lot of research these days is not reproducible. It is something that should definitely be avoided, the going back and saying, *OK, so I didn't find it first go around, so I am going to change things around and try again.*

**Suzanne:** So, the interesting thing about that is that one of the other things from...I'll call it classical scientific method, is to say, refine your hypothesis based on the data that you get. And so what you're saying is, *Be cautious about that*, because data, what we've learned about data is not all data is equivalent. Not all data has the same relevance to the question that we're asking. And so, not only do we need to change and refine the hypothesis, but we need to understand

what are the anomalies in the data so we are not inadvertently using anomalies as a way of focusing a refinement. Is that fair?

**Leigh:** In cybersecurity, people often see an anomaly and assume that means maliciousness. It's called the fallacy of anomalies. It's, *Oh, I found the weird thing in here, so therefore it must be bad*.

**Jonathan:** Please don't do that. Please, everyone stop doing that.

**Leigh:** Yes.

**Suzanne:** *An anomaly does not equal bad* is one of the takeaways from this podcast.

**Leigh:** An anomaly does not equal bad.

**Jonathan:** Well, unless you define your security policy, that everything that happens less than two percent of the time is not allowed by the security policy, in which case then, sure.

**Leigh:** The other thing about data is, and the Internet is, everyone has their own view of the Internet. So, if I collect a certain kind of say, DNS data, and someone else also collects DNS data, we are not going to get the same exact dataset. It's influenced by our location, by what we do…

**Suzanne:** By the dynamism of the infrastructure that we are dealing with.

**Leigh:** Yes, the distributed nature. It's the same with DNS, the same with routing, malware blocklist. Jono and I showed blocklists, it is definitely true. So, it is all in how you collect the data. So, by designing an experiment for your weirdnesses of data, you are not making it useful for someone else.

**Suzanne:** I am just thinking of applying this to datasets that I have worked with in the past, not from a security viewpoint, but just from answering other questions. How do you understand which anomalous, unexpected data to pay attention to for a refinement purpose and which to ignore? Is there any kind of a rubric to help people with that, or am I getting way too deep into this?

**Leigh:** It is not a subject-matter-expert thing, for one thing. On the other hand, it is a very difficult question. It is why we have verification and validation. It is why you save part of your data aside and see if you can replicate it in this, you know, validate it. It is why we have corroboration. Jono and I did a report, it is 300 pages on blacklists. It actually was corroborated soon after we released it by someone else using a completely different dataset. I was told he

started his talk with, *They are wrong. They are completely wrong, and I am going to show they are wrong*, and at the very end, he was like, *No, they are right.*

**Suzanne:** That must have been a fun day. Jono, did you have something to add to that?

**Jonathan:** Yes, I mean, I think there are a number of different desirable properties of a study, whether it is software engineering or cybersecurity or psychology or astronomy or whatever. Different fields often use different terms. But, I think that there is an umbrella around consistency, generalizability, transparency, and containment of harms. So, if you have anomalies in your dataset, one of the things that you are looking at is a lack of consistency. And, there are a lot of general categories of ways that can happen. You can have *your study was designed in a way that was inconsistent*. You can have *your tool is constructed, engineered, in a way that is inconsistent*. If there is a scratch on your telescope, you are going to have consistent data-collection errors, but there is a pretty clear mechanism for how that happens, and there is a way to fix it, if it matters. There are less obvious things with psychology and study design and stuff like that that cause consistent data-manipulation errors. Finding artifacts like that, which are not always super obvious, that might be part of the study design, so part of an inconsistency thing, are difficult for sure, but mostly requires you to go looking at what other people have done. If you are designing a new tool, studying the consistency of that tool on known systems where you know the answer, is just a totally different project than using that tool to go learn something new about something you don't understand. I think one of the things we see in cybersecurity is people don't separate those things.

**Suzanne:** To me this has an analogy to some of the techniques that we use for modeling and simulation. When we are looking at a simulation, we look at a known dataset to validate the simulation against, because we know what the answer is. We know what the actual physical space looks like, etc., etc., and we use that. So, that technique is one that is very valid to also use in applying these methods to cybersecurity and data analysis in particular is what I am hearing from you. So, other things about data collection and use that come out of scientific practices and methods that cybersecurity as a field may not be paying enough attention to right now. You mentioned one, which is the assumption that *the data I know about and the data I don't know about are going to be equivalent*. What are some others that you can think of?

**Jonathan:** My favorite page on the whole Internet is Wikipedia's list of cognitive biases.

**Suzanne:** Ah, yes.

**Jonathan:** Essentially, that list is all of the documented ways in which human brains do not behave according to the statistical rules that you would like them to if you were doing completely rational and proper statistical inference. For almost all of them, there are very good

evolutionary reasons why your brain makes those shortcuts, but also there are the 200 ways in which you are going to trick yourself, when you think that you know what you are doing in a scientific study.

**Suzanne:** So, going back to that basic set of cognitive biases and applying it to the realm of cybersecurity. I think, Leigh, in some ways, you brought one up—I'm not sure if it was Jon or Leigh—that brought up the one that is, *Any anomaly is bad*, is essentially a cognitive-bias issue of *I am biased towards anything that is unexpected, that the cause of that is some kind of malware, some kind of a bad event.* So, that is an example.

**Jonathan:** Sure.

**Suzanne:** I know there are lots of other examples, and we will be sure that the link to that page is in the transcript. I also have a friend who has a poster of something similar in his office. So that is why I smiled when you brought that up. It's like, *Oh, yeah, there are a whole bunch of these that we can deal with*. That is a good reminder, that this is...and I take your point earlier about the cyber part of this is actually, in some ways, it is the socio aspect to it. I mean, you could almost call cybersecurity socio-security, except it is harder to say. When you get into the social sciences, you get into a much, much bigger realm of cognitive bias and difficulty and subjective data vs. objective and all of those kinds of things. But what I am hearing from you is, that distinction is important, that we are not just dealing with the technical aspects of security. We have got to look at that larger space, from a scientific viewpoint.

**Jonathan:** Well, I think that is an interesting point, Suze, because there are actually two things that you are bringing up there. One is like behavioral economics of studying how people actually behave, given all these cognitive biases, which is super relevant. But even if you were studying a purely physical object, like if you are doing astronomy, no humans involved. You are a human. You have cognitive biases. You still need to be very careful when designing studies and analyzing results that you are not falling into these traps. And so, even if we have listeners who think like, *I just study kernel development and cryptography algorithms. I don't need to worry about this because I don't touch human beings*. You are a human. You still actually do need to go there more. Be aware of this so that when you design studies, you are not tricking yourself.

**Suzanne:** OK. Is that one of the functions of a journal like DTRAP where people publish? I want to encourage people that have insights and that have viewpoints they want to get validation for… I can imagine that outside the SEI… I know [at the SEI] we have a lot of mechanisms for getting verification and validation in our research; but outside of that realm, you may not have very many avenues for getting that validation. Is that one of the things that DTRAP has actually been able to help with?

**Leigh:** I think so because we created a paper type called a *field note*. We called it a field note because it came from high-energy physics, and in high-energy physics, when you come up with a new particle, you don't publish an entire new paper about the new particle. There is a short paper that describes what the particle is. And so, a field note could be someone saying, *Hey, I found this new thing, and this is what I tried with it. Is it working?* We created it for a couple of reasons. We created it because we can't actually ask people from industry to write full academic papers. I have been in industry. If I told my boss I was writing a 25-page academic paper, he would probably still be laughing, but I could say I am writing a shorter paper, and he'd say, *Fine, go to it*. So we created it for that, and we also created it because we want to see the new ideas. We want to see the new things that are coming out, the new particles, the new ideas. We think it is working. It has been a little rough getting going, because most people want to go*, Oh, I'm reviewing an academic paper*, and we're like, *No, there's a little bit different here. It's not as in-depth of a review.* We also have two columns. We have one that's called, "With the Benefit of Hindsight," which we call, hey, what did you learn from the past? And the other is, "From Research to Practice." The goal from that was how do you take research and turn it into practice? And what did you learn?

**Suzanne:** Gotcha. OK, so we have got some mechanisms that we have talked about for thinking about scientific practice and methods being applied. We are going beyond falsification as the only way that we generate hypotheses. We have looked at some common pitfalls that we can get into, cognitive bias being my favorite. I am with you on that one, Jono. Many of the things that we're talking about are in your book, *Using Science in Cybersecurity*. What is next? What is the sequel to this book? What are some of the topics that we've talked about today or haven't talked about today that you think need to be addressed and that you would like to be able to cover. What still needs to be done in this area?

**Leigh:** Well, I am actually working on somewhat of a sequel now with two authors.

**Suzanne:** I'm so surprised to hear you say that.

**Leigh:** We are looking more at the pitfalls in cybersecurity. We are taking a closer look, and it's not just for researchers. It's not really just for people at the low level of *I'm trying to deal with the sense of it.* We want everyone who has to deal with cybersecurity in some way, to understand there's some fallacies in thinking. There are cognitive fallacies. There are bad assumptions that people make. There are misunderstandings about vulnerabilities and malware. So, that is my next step, moving forward with this.

**Suzanne:** OK, so getting a better understanding of where some of the pitfalls are, especially, I'm guessing, the ones that wouldn't be obvious to somebody like me, who wants to be secure in the way that I act with the world, but may not be as aware of some of the nuances, of pitfalls, that I

may be bringing into my own practice and trying to be a secure operator within the areas that I'm playing in.

**Leigh:** Yes. Yes.

**Suzanne:** OK. What about you, Jono? What are some topics that would be in your sequel?

**Jonathan:** Yeah. So, Suze, I found on page 37 a typo. I would like to fix that, as I was reading about consistency, to answer your last question.

**Suzanne:** I have written a book. I am there with you. Mine is on page 17. No matter how many times you read it, it's almost like it's required.

**Jonathan:** I'll give a high five to whoever figures out which typo that is first and sends me an email. No, so I am partly involved in some stuff at FIRST, the Forum of Incident Response and Security Teams. I have also been involved in an effort around ethics FIRST, so like ethics for incident response and security teams. One of the duties there that is being proposed is a duty for evidence-based reasoning. So, in order to ethically conduct incident response and be a security team, you have to provide evidence, work from evidence, even if there is stuff that for good reason you can't share, you need to have done it in a way that is consistent with reliable evidence gathering, consistency, appropriate generalizability for what claims you are making. I also am on the Common Vulnerability Scoring System [CVSS] SIG [special interest group], a standard for how to score the technical severity of vulnerabilities. We need to get these sorts of evidence-collection analyses in, just all of these basically scientific methods, into our standard processes. Like right now, CVSS—and I've written about this enough—but CVSS is not transparent in how it prioritizes things.

**Suzanne:** So there may be context in which the prioritization is not as applicable, if you don't know what the criteria were for setting that prioritization.

**Jonathan:** Yes, and I have good reason to believe that the criteria have some other problems. There was like 100 people in a room saying which things were the worst. Then someone set a curve to it, and they didn't really explain how or why they did that or whatever. I think that that is just one example of one that I am familiar with. I think there are a lot of standards which are compliance checklists without a whole lot of evidence for the efficacy of what is on the checklist. The places that I want to go with this are getting these sorts of scientific best practices into operations. If that is through standards or through I don't know what insurance companies are expecting if they are going to give you cybersecurity insurance or through norms of industry groups, or…I guess FIRST is not an industry group, because it is a lot of non-profits and government organizations. But, there are other industry groups that I think you could also…

**Suzanne:** Within the community as a whole?

**Jonathan:** Just generally within the community. People actually...people understanding this better.

**Suzanne:** What I heard you say earlier, you are looking for this evolution to occur, not just related to the technical aspects, but also to the sociotechnical…

**Jonathan:** Yes that is right.

**Suzanne:** …and helping people, because the ethics that you brought to mind, the idea that there is an education process that is needed for people to understand where cognitive bias is related to ethics, for example. There may be particular types of cognitive bias that are more prevalent when you are dealing with ethics-related questions than when you are dealing with purely technical questions.

**Jonathan:** Yes, but I think it is super hard because, at some point, we are talking about…Everyone is working from the scientific method, singular, that they were taught in middle school, whatever, 8th grade, 7th grade, whatever grade, right? Maybe we need to start looking at making that curriculum more nuanced, so that we are, but I don't have any…

**Suzanne:** So you need to get into the STEM committee because that is really what you are talking about, is modifying the STEM curricula to be more inclusive of different aspects of scientific methods. Start them when they are young. Yep, I see that.

All right. Any topics that we didn't cover that you guys wanted to talk about today, we've been around a whole bunch of different ideas, and I for one have been enjoying this conversation a lot. But, anything that you wanted to make sure that our viewers know about, that we didn't cover?

**Leigh:** One of my favorite topics in this research is that negative results are still results. People hear the word negative, and they want to toss it, because it said *negative*. So therefore, *bad*. Negative results are not bad results. Negative results are just, *I didn't prove my hypothesis*. Cybersecurity is an ever-changing field. When I started in this business, what I was concerned with is completely different from what we are concerned with now. I have never dealt with a ransomware attack, when I was dealing with it, and my first instance of malware was dealing with something that came on a floppy.

**Suzanne:** Ten-inch or three-and-a-half?

**Leigh:** Five-and-a-quarter.

**Suzanne:** Five-and-a-quarter. Okay. I started with the 10.

**Leigh:** The negative results essentially can be used to say, *Hey, what your assumptions were on how cybersecurity worked five years ago are no longer valid*. And you need…

**Jonathan:** Look. No, no, no.

**Suzanne:** Oh, my goodness.

**Leigh:** I hit the wrong topic with Jono.

**Jonathan:** This is like I, I think I've got a solitary war against…

**Leigh:** Yes, you do.

**Jonathan:** …the perception of these terms. I talked about moving past falsification, because it's not super helpful. The whole idea of positive and negative results requires an A-B randomized control trial that is appropriately controlled, and that is only one kind of study. So in the first place, it is often getting applied to like a case study. We can't have negative results in a case study, because you're not having a test, really, in the right way. But a better thing to call negative results would be unsurprising results. And so, what Leigh was saying was, *Oh, maybe you learned something, that something you thought was true, like now is not true*. That is not a negative result. That may be surprising. Just because you didn't know it, doesn't mean it's a negative result. In general, it means that you have set up some particular null-hypothesis test, which is a very specific statistical thing, and you failed to reject the null hypothesis, which just means that there is no evidence that what you were testing for actually matters. But that might be surprising.

**Suzanne:** Fair enough. But I think, Leigh, to go back to what you were talking about, with the evolution of that knowledge of...I would reframe it, to say the evolution of the knowledge from what was surprising to now, what is not surprising. That changes based on context. Is that fair to say?

**Leigh:** Yes.

**Jonathan:** I think with DTRAP and some of these other efforts, there is this big difference between a negative result, where you have done everything very well, and the result is not surprising, and a non-result, where you have done a crap job of designing the experiment, and so you don't get any information. I do not want to publish non-results unless you are telling me how not to design experiments. I want to publish all of the negative results, because it tells me what I don't have to repeat.

**Suzanne:** OK.

**Leigh:** Yes. So, maybe we should couch it as *surprising* and *non-surprising*, but yes, I want to know what I don't have to repeat. It is not just cybersecurity where people don't report these results. There have been studies that show that these results get hidden, or people use data dredging, which I mentioned before, to basically go back and reframe the results, so they get a quote "positive result," and they can publish, because journals don't like…

**Suzanne:** They don't want to hear that there was no surprise here.

**Leigh:** Yes. And I'm actually looking at…My term as editor-In-chief of the current...of DTRAP ends in December, with ACM. They have a term limit. My next journal is entirely possibly going to be on this topic.

**Suzanne:** OK. So I was going to suggest a column, a new column, "Things That Were Not Surprising."

**Leigh:** No, actually, I have interest from enough people that I believe I am going to start another journal.

**Suzanne:** Fair enough. All right. With that, I am going to go ahead and close today's podcast. I think we have lots of things to talk about in the future as this work evolves. And I really, really appreciate this discussion, because I am one of the people that is much more in tune and forward in terms of the way I think about science, and the fact that we are applying that to this very complex sociotechnical set of problems, that is very encouraging to me. So, I want to thank both of you for this conversation and for your work and your book, and I want to thank you for joining us both today. To our viewers, as always, we will include links in our transcripts to all the resources we mentioned: Wikipedia, DTRAP, the workshops, all that stuff. You will get references in the transcript, and I want to thank you, once again, for joining us.

**Jonathan:** Thank you, Suze. Thank you, Leigh.

**Leigh:** Thank you, Jono, and thank you, Suze.

**Suzanne:** You are very, very welcome.

*Thanks for joining us. This episode is available where you download podcasts, including SoundCloud, Stitcher, TuneIn Radio, Google Podcasts, and Apple Podcasts. It is also available on the SEI website at sei.cmu.edu/podcasts and the SEI's YouTube channel. This copyrighted work is made available through the Software Engineering Institute, a federally funded research and development center sponsored by the U.S. Department of Defense. For more information about the SEI and this work, please visit www.sei.cmu.edu. As always, if you have any questions, please do not hesitate to email us at info@sei.cmu.edu. Thank you.*