On the Science of Security

John D. McLean | Naval Research Laboratory

In a recent article,[1] Cormac Herley and P.C. van Oorschot present an extremely informed discussion on the philosophy of science in general and the prospects for a science of security in particular. [Editor's note: For related work, see the sidebar.] Although I agree with most of what Herley and van Oorschot say, they include a claim—made originally by Herley and van Oorschot in "SOK: Science, Security, and the Elusive Goal of Security as a Scientific Pursuit"[2] and by Herley in an even earlier paper[3]—which I think is misguided and which has the danger of doing damage to the field. All three articles claim that computer security fails to avoid unfalsifiable claims and statements. Insofar as this statement is simply pointing out that security practitioners often make statements that are vague or imprecise, I don't disagree. One could argue that we all know what they really mean, but as Herley and van Oorschot point out in "SOK: Science, Security, and the Elusive Goal of Security as a Scientific Pursuit,"[2] I've already noted that the use of hidden assumptions is the path to neither science nor security.[4] However, Herley, in his earlier paper, and Herley and van Oorschot, in their more recent articles, clearly believe that this unfalsifiability is somehow inherent in the study of security, per se. As the authors put it in both of their joint articles: "claims of necessary conditions for real-world security are unfalsifiable. Claims of necessary conditions for formally-defined security are tautological restatements of assumptions." To gain a better understanding of why I think that this claim is misguided, it is worthwhile to consider an example that is presented in both articles in support of it.

Herley and van Oorschot note that the claim that a password must have certain properties to be secure is unfalsifiable. For ease of exposition, let's assume that the property in question is the property of containing, at least, 8 symbols drawn from an 80 symbol character set. The reason that this claim is held to be unfalsifiable is that to disprove the claim one would need to find a secure password that does not satisfy those properties. But how could anyone possibly demonstrate that any password is secure? One could point out that nobody has yet broken the password, but a skeptic could always reply: "Not yet, but keep waiting." As Herley and van Oorschot put it: "we can't find a secure password because there is also no way to show that a password is safe against not-yet-known attacks,"[1] and "successfully avoiding harm in the past is no guarantee about the future."[2] If correct, the argument demonstrates that computer security can never be truly scientific and that the computer security community is condemned to forever play "penetrate and patch" with adversaries who have shown a remarkable ability to find weaknesses in our systems. Fortunately for the science and practice of computer security, the argument rests on a misunderstanding both of computer security and of science. The nature of this misunderstanding illuminates both.

Consider Newton's theory of gravity, a theory that Herley and van Oorschot cite and a theory that is often regarded as being an exemplar of good science. An analogue of the claim that a secure password contains more than 8 characters is the claim that a metal ball released near the surface of the earth will fall. Just as a lack of precision makes the password claim unfalsifiable, it makes this claim about gravity unfalsifiable as well. I'm not referring to the fact that auxiliary assumptions are required—for example, the assumption that the ball is not being interfered with by a magnetic field—and that these assumptions can always be questioned if the predicted event does not come to pass. Auxiliary assumptions are always required and can always be jettisoned in the face of adverse evidence. My point is simpler. If the ball doesn't fall, a skeptic can, as in the password case, say: "Not yet, but keep waiting." Although physicists make such statements, especially when explaining physics to lay audiences, these statements do not have any bearing on the question of whether physics is a science. What makes a theory of gravity scientific is that it yields predictions that are falsifiable, for example, the prediction that a ball dropped in a vacuum from a height of 20 meters will hit the ground in 2.02 seconds.

I would argue that exactly analogous statements can be made

Copublished by the IEEE Computer and Reliability Societies

in the field of computer security. For example, if I were fortunate enough to own a supercomputer that can compute a SHA512 hash in 1.7×10^{-11} seconds, I can then make the prediction that if I have a hashed password comprising 8 symbols chosen from an 80-symbol character set, my computer can crack the password in under 8 hours. It is the falsifiability of this prediction that makes computer security a science.

At this point, Herley and van Oorschot may reply that the statement I just held up as being scientific is merely tautological. I would agree that the statement that a computer that can search the entire space of 8-symbol passwords chosen from an 80-symbol character set in 8 hours can find a given 8-symbol password from that character set in 8 hours is a tautology. However, the claim in question is that a specific computer can break the password, and there's nothing tautological about that claim. It may be a mathematical truth that a computer that can compute a single SHA512 hash in 1.7×10^{-11} seconds can compute 80^8 hashes in 8 hours, but it's an empirical, not mathematical, truth that my computer can. A similar point holds for the theory of gravity. Although it's a mathematical truth that a ball dropped from a height of 20 meters with an acceleration of 9.75 meters/second^2 will hit the ground in 2.02 seconds, it's not a mathematical truth that a ball dropped from a height of 20 meters on the surface of this planet will hit the ground in 2.02 seconds. The fact that the acceleration due to gravity near the earth's surface is 9.75 meters/second^2 is empirical, but so is the fact that a given computer can compute a SHA512 hash in 1.7×10^{-11} seconds. The first fact depends on the mass of the earth, the value of g, and so forth. The second fact depends on various properties of specific circuits, the architecture of the computer, the existence of certain algorithms, and so forth. It's also worth noting that the claim about finding the password remains an empirical truth whether or not I can actually manage to obtain the hashed password, just as my claim about the behavior of a dropped ball remains an empirical truth whether or not I actually manage to create the required vacuum.

An obvious response at this point is to state that even if there is nothing inherently unscientific about computer security, the same cannot be said for security engineering. Herley and van Oorschot could point out that while it's all well and good to state that the Science of Security community should be clear as to what we can scientifically state and what are mere rules of thumb derived from that science—for instance, avoid 8-character passwords—the fact remains that when it actually comes to engineering secure systems, we are not nearly as far along as the physical sciences are with respect to engineering physical structures. As they say in "Science of Security: Combining Theory and Measurement to Reflect the Observable,"[1] "[T]he real-world system in which attacks actually occur simply contains threat vectors beyond those considered in the abstract system of deductive reasoning." Although we know that certain properties are necessary to render passwords secure from certain classes of attack, we don't know what properties would make a password secure from all attacks. In other words, we know some scientific facts about passwords, but this knowledge is insufficient to engineer completely secure passwords. However, I'd point out that this is not because security is somehow inherently unscientific or that security engineering is inherently impossible, it's just that the science of security is currently not sufficiently developed to support security engineering as adequately as we would like.

That said, it's worthwhile emphasizing that the difference between security engineering and, for example, structural engineering is a matter of degree rather than kind. We may not know how to build a system that is secure from all attacks, but we don't know how to construct a bridge that is indestructible either. Instead, we use physical science to make bridges that can survive exposure to certain classes of natural events. If a catastrophe occurs outside that class—for instance, being hit by a meteor the size of Manhattan—all bets are off. The same is true for computer security. It's just that in the case of security, we have an intelligent adversary trying to find and exploit flaws, rather than an indifferent, if not quite benign, natural world that is, to some extent, predictable.

To put security engineering in perspective, it's worthwhile remembering that computer security is, as is computer science, a very young field. The term *Physics* first appeared in print approximately 2,400 years ago, and it took over 2,000 years to advance from Aristotle to Newton, another 175 for Maxwell's equations, another 50 for General Relativity, another 10 for Schrodinger's wave equations, and we've had 90 years to build on that. By contrast, as Herley and van Oorschot point out, the principles of computer security weren't set out until 1975. Although we admittedly have a long way to go in developing a sufficient body of science to support secure engineering that is as resilient as we would like, I think the field has made remarkable progress since its inception. Claiming that its empirical claims are somehow inherently unscientific is incorrect and a disservice. ■

Response to "On the Science of Security"

Cormac Herley | Microsoft Research

P.C. Van Oorschot | Carleton University

We thank Dr. McLean for his interest in our work, and the editor for the opportunity to respond.

We confirm the statements made in the articles[1–3] cited in McLean's letter: claims of necessary conditions for security (in the sense of avoidance of future harm) are unfalsifiable. There is nothing mysterious about this statement. To falsify the claim "To be secure, you must do X," requires pointing to something secure that doesn't do X. The difficulty with that is convincingly showing that the "something" is secure: How do you convincingly show that a real system is secure? To show insecurity is easy (demonstrate an attack); to demonstrate security is not (we can't "observe" that something is immune to attack). A main difficulty comes from vague and misleading use of words such as "secure" and "security proof." It's best if we avoid language ambiguity, a problem long-known in many fields of science. We argue that rather than making assertions that a system is "secure," the community is better served by statements about resilience to specific known attacks, and about specific observable outcomes.

Of course you can instead choose to define security to be possession of certain properties. Yes, you can empirically verify that a machine computes 80^8 hashes in 8 hours—but that does not imply a necessary condition for avoidance of future harm. We can define security in many different ways, but a great deal of confusion is generated by moving between definitions; that confusion can be avoided by avoiding use of terms that others misinterpret.

We are puzzled as to how McLean reaches the conclusion that this "demonstrates that computer security can never be truly scientific." If claims of a certain type are unfalsifiable, as we have pointed out, then the constructive approach, and one we advocate, is to concentrate on different types of claims, rather than concluding as McLean does that the "community is condemned to forever play 'penetrate and patch.'" Examples are claims of improvement in average outcomes over finite timeframes. We advocate for the use of scientific approaches where they help—as well as engineering and other approaches, wherever those also help.

Finally, we are more optimistic than McLean on the question of damage to the field. Science is about open discussion. We hope that no one will accept any of our claims simply on authority, and instead will examine them critically and make up their own minds. Our articles are written for a wide audience, with "Science of Security: Combining Theory and Measurement to Reflect the Observable,"[1] having been solicited as a shorter version of "SOK: Science, Security, and the Elusive Goal of Security as a Scientific Pursuit,"[2] and to stimulate further discussion and thought. We would be delighted if more people take the time to read them, and perhaps McLean's comments may have this positive effect. ■

References

1. C. Herley and P.C. van Oorschot, "Science of Security: Combining Theory and Measurement to Reflect the Observable," *IEEE Security & Privacy*, January/February 2018, pp. 12–22.
2. C. Herley and P.C. van Oorschot, "SOK: Science, Security, and the Elusive Goal of Security as a Scientific Pursuit," *Proc. 2017 IEEE Symp. on Security and Privacy*, pp. 99–120.
3. C. Herley, "Unfalsifiability of Security Claims," *Proceedings of the National Academy of Sciences of the USA*, vol. 13, no. 23, 2016, pp. 6415–6420.
4. J. McLean, "Reasoning about Security Models," *Proc. IEEE Symp. Security and Privacy*, 1987, pp. 123–133.
5. C.E. Landwehr, "History of US Government Investments in Cybersecurity Research: A Personal Perspective," *Proc. 2010 IEEE Symposium on Security & Privacy*, 2010, pp. 14–20.
6. J. McLean, "On the Science of Security," *IEEE Security & Privacy*, vol. 16, no. 3, 2018, pp. 6–7.
7. C. Herley and P.C. van Oorschot, "Science of Security: Combining

NSA's Research Initiative in the Science of Cybersecurity

Research into foundational solutions for problems in computer security, data security, information technology security, and now cybersecurity has been underway for nearly a half century, supported by funding from numerous research agencies.[5] Yet an explicit focus on developing a *science* of cybersecurity, discussed by Herley and van Oorschot and is relatively recent.[6,7] The primary US initiative in this direction is the Science of Security program administered by the National Security Agency. This program was stimulated by concerns among government research funders that triggered the 2008 workshop and 2010 JASON study cited by Herley and van Oorschot.[6]

In 2008, the US government initiated a Comprehensive National Cybersecurity Initiative (CNCI).[8] It included a number of programs targeted at reducing vulnerabilities and dealing with cyberattacks.[9] Most of these improved operational systems, but there were efforts to advance education and research as well. Growing interest in establishing scientific foundations for cybersecurity was reflected in a focus issue of this magazine in May/June 2011.[10] In December 2011, the National Information Technology Research and Development (NITRD) Program published a strategic plan for federal cybersecurity research and development calling for thrusts in inducing change, accelerating transition to practice, maximizing research impact and, notably, developing scientific foundations.[11] Research funding allocated under the CNCI's research effort enabled NSA's Science of Security (SoS) research initiative.

Two issues of NSA's research magazine, *The Next Wave*, reflect the initiation of the program[12] and describe related, coordinated initiatives funded by the US Army Research Laboratory, the Air Force Office of Scientific Research, the US National Science Foundation, and the UK Government Communications Headquarters.[13]

NSA's SoS initiative includes three main thrusts: engage the academic community for foundational research, promote rigorous scientific principles in research, and grow the SoS research community. Building on a platform developed through National Science Foundation funding, NSA created a public SoS "virtual organization" (https://cps-vo.org/group/SoS) to enable communication and to document the initiative's activities. To date, 1,400 members have joined.

Research has been funded directly through "lablets"—university-based research centers. An initial three centers grew to four, competitively selected, and is expected soon to grow to six. Each of these is charged not only to conduct SoS research, but to collaborate with other institutions as a means of growing the community. The results to date include more than 550 publications with authors from about 175 institutions. With the assistance of the lablet researchers, an annual Hot Topics in Science of Security (HotSoS) conference has been created, also as a way of building a community.

The SoS initiative also instituted a "best paper" award, to be given to the paper published in the preceding calendar year deemed to have done the most to advance the science of cybersecurity by building scientific foundations or exemplifying rigorous scientific methods. A set of distinguished experts provide individual comments to the NSA decision maker in this effort, now in its sixth year. NSA has also instituted a program of awards given at K–12 science fairs to students with projects relating to the science of cybersecurity.

The creators of the SoS program noticed that other fields of science often advance by focusing their efforts on the study of particular hard problems or testing of specific theories. Again with the assistance of lablet researchers, the program identified a set of five hard problems—resilient architectures, metrics, scalability and composability, secure collaboration, and understanding and accounting for human behavior—and has publicized them in an attempt to focus and advance the field.[14]

The dialog in this article regarding the status of cybersecurity as a science and the possibilities for making the field more scientific is very much consistent with the thrust of these research initiatives. *IEEE Security & Privacy* magazine welcomes further correspondence and contributions on the topic.

Theory and Measurement to Reflect the Observable," *IEEE Security & Privacy*, vol. 16, no. 1, 2018, pp. 12–23.

8. G.W. Bush, *Cybersecurity Policy. National Security Presidential Directive 54*, 8 Jan. 2008, pp. 10–15; https://fas.org/irp/offdocs/nspd/nspd-54.pdf.
9. O.S. Saydjari, "Launching into the Cyberspace Race: An Interview with Melissa E. Hathaway," *IEEE Security & Privacy*, vol. 6, no. 6, 2008, pp. 11–17.
10. D. Evans and S. Stolfo, "Guest Editors' Introduction: Science of Security," *IEEE Security & Privacy*, vol. 9, no. 3, 2011, pp. 18–19.

11. *Trustworthy Cyberspace: Strategic Plan for the Federal Cybersecurity Research and Development Program,* Executive Office of the President, National Science and Technology Council, 2011; https://www.nitrd.gov/pubs/Fed_Cybersecurity_RD_Strategic_Plan_2011.pdf.
12. R.V. Meushaw, ed., "Developing a Blueprint for a Science of Cybersecurity," *The Next Wave,* vol. 19, no. 2, 2012; https://www.nsa.gov/resources/everyone/digital-media-center/publications/the-next-wave/assets/files/TNW-19-2.pdf.
13. F.R. Chang, ed., "Building a National Program for Cybersecurity Science," *The Next Wave,* vol. 19, no. 4, 2012; https://www.nsa.gov/resources/everyone/digital-media-center/publications/the-next-wave/assets/files/TNW-19-4.pdf.
14. D.M. Nicol et al., "Science of Security Hard Problems: A Lablet Perspective," 2012; https://cps-vo.org/node/6394.