Practicing a Science of Security

A Philosophy of Science Perspective

Jonathan M. Spring
University College London
Gower Street
London WC1E 6BT
jspring@cs.ucl.ac.uk

Tyler Moore
The University of Tulsa
800 South Tucker Drive
Tulsa, OK 74104-9700
tyler-moore@utulsa.edu

David Pym
University College London
London WC1E 6BT
Alan Turing Institute
d.pym@ucl.ac.uk

ABSTRACT

Our goal is to refocus the question about cybersecurity research from 'is this process scientific' to 'why is this scientific process producing unsatisfactory results'. We focus on five common complaints that claim cybersecurity is not or cannot be scientific. Many of these complaints presume views associated with the philosophical school known as Logical Empiricism that more recent scholarship has largely modified or rejected. Modern philosophy of science, supported by mathematical modeling methods, provides constructive resources to mitigate all purported challenges to a science of security. Therefore, we argue the community currently practices a science of cybersecurity. A philosophy of science perspective suggests the following form of practice: *structured observation to seek intelligible explanations of phenomena, evaluating explanations in many ways, with specialized fields (including engineering and forensics) constraining explanations within their own expertise, intertranslating where necessary.* A natural question to pursue in future work is how collecting, evaluating, and analyzing evidence for such explanations is different in security than other sciences.

KEYWORDS

security research; science of security; cybersecurity; history of science; philosophy of science; ethics of security

ACM Reference format:
Jonathan M. Spring, Tyler Moore, and David Pym. 2017. Practicing a Science of Security. In *Proceedings of 2017 New Security Paradigms Workshop, Santa Cruz, CA, USA, October 1–4, 2017 (NSPW 2017),* 18 pages.
DOI: 10.1145/3171533.3171540

1 INTRODUCTION

There has been a prominent call to improve the research and practice of information security by making it more 'scientific'. Its proponents claim a science of security is needed for ongoing progress. Per the historian of science Dear, scientific is used here as "a very prestigious label that we apply to those bodies of knowledge reckoned to be most solidly grounded in evidence, critical experimentation and observation, and rigorous reasoning" [20, p. 1]. We take our definition of security from RFC 4949: "measures taken to protect a system" [81]; of course see the RFC for the meaning of measure, protect, and system. As a starting point, then, we consider a science of security to be the label we should apply to the most solidly grounded bodies of knowledge about measures one can take to protect a system.

The following items have resonated as serious obstacles to the practice of a science of security:

- Experiments are impossible in practice, because they are unethical or too risky;
- Reproducibility is impossible;
- There are no laws of nature in security;
- Information security will not be a science until we all agree on a common language or ontology of terms;
- Computer science is 'just engineering' and not science at all: questions of science of security are misplaced.

We will argue that a philosophy of science perspective shows these obstacles are either misguided or can be overcome. The purported obstacles are frequently not genuine challenges because they rely on outdated conceptions of science, which yields a simplistic idea of evaluating evidence for claims (falsification) and a naïve reductionism to universal laws that supposedly underpin all scientific endeavors. Alternatively, modern philosophy of science tends to describe, if applied adequately to security, what good security practitioners already do. Security is, as practiced, already a kind of science. Table 1 summarizes our positive perspective on executing a science of security.

Section 2 provides a brief background on the logical empiricist movement within philosophy of science. Section 3 examines prior statements to detail the obstacles to practicing a science of security. Section 4 explains how philosophy of science informs the scientific

Experiments are untenable	Structured observations of the empirical world
Reproducibility is impossible	Evaluate by repetition, replication, variation, reproduction, and/or corroboration
No laws of nature	Mechanistic explanation of phenomena to make nature intelligible
No single ontology	Specialization necessitates translation
'Just' engineering	Both science and engineering are necessary

Table 1: Five common complaints raised by the science of cybersecurity community and positive reframing from the philosophy of science literature.

NSPW 2017, Santa Cruz, CA, USA

DOI: 10.1145/3171533.3171540

process already taking place in cybersecurity research, and Section 5 suggests some constructive steps forward for improving the reliability and growth of general, sharable knowledge in security.

2 PHILOSOPHY OF SCIENCE PRIMER

Philosophy of science is a field that has developed as a discourse on top of science: a second-order reflection upon the first-order operation of the sciences [89]. For three centuries, the scholars we now recognize as scientists were called 'natural philosophers', and there was no separate group of philosophers of science. In inter-war Vienna, a group of thinkers who identified as 'the Vienna Circle' came to challenge both the prevailing metaphysics and political Romanticism (i.e., the Church and European facism).[1] This movement emphasized themes of observation of the world, trust in science, high value on math and logic, and modernism. A key movement of the Circle has come to be called *logical empiricism*, for its reliance on logical rules based on empirical observations.[2]

We briefly introduce two of the main tenets of logical empiricism: (i) empiricism and verification and (ii) unity or reduction of scientific fields [15]. These tenets coalesced in the 1930s, were refined through the 50s, and by 1970 had suffered ample critiques to be changed beyond recognition. This historical trajectory makes it intellectually dangerous to rely upon logical empiricist arguments or concepts uncritically. Yet, Section 3 finds much logical-empiricist work uncritically assimilated in current statements on science of cybersecurity.

Empiricism and verification. Statements testable by observation were considered to be the only "cognitively meaningful" statements [89]. Although logic and mathematics are the most reliable forms of reasoning, logical empiricists did not take them to rely on observation but instead accepted them as true by definition, following Russell and early Wittgenstein. Therefore, according to the logical empiricist view, the key scientific challenges are how to verify a statement is in fact about the world, and how to meaningfully integrate observations into logic and mathematics. Such integration is necessary for science to be useful. Integrating observations into deductive logical statements is also a response to Hume, two centuries earlier, and his famous problem of induction. Hume, in broad strokes, argues that no matter how many times we observe the sun to rise, we cannot prove (in the sense of deductive proof) that the sun will rise tomorrow based on the observations.

Consistent with logical empiricism, Carnap proposed a method for *verification* by working on atomic elements of logical sentences, and expanding observational sentences out based on rules from atomic observations [15]. The goal of empiricism is to be grounded in observations. The goal of verification is to integrate those observations into a framework of general knowledge, in the form of statements in first-order logic, that can justify predictions. Carnap thus links induction and deduction, bypassing Hume's complaint.

Yet it became clear that verification might not always be achievable. It is against this backdrop that Popper proposed the more limited objective of falsification [74], which claims we cannot verify logical statements at all. Instead, Popper asserts that the best we can do is hope to falsify them.[3]

Even the more limited goal of falsification was shown to be untenable with Kuhn's challenge to Popper in 1962 [54]. Kuhn refutes the premise that scientists operate on logical statements. Rather, he argues that key examples, literally 'paradigms', are scientists' operative cognitive model. Later work in philosophy of science has refined the shape of these cognitive models —- one prominent method is as *mechanistic explanations* (see, e.g., [37]) — and improved understanding of how data are processed to provide evidence for phenomena (see, e.g., [7]).

Even ignoring Kuhn's socio-scientific critique, falsification is about mapping observations into logic. Popper is silent on designing reliable observations and choosing what logic or conceptual framework in which we should reason. These two problems are more important, and would provide more actionable advice, than whether something is falsifiable. More useful than falsification are modern discussions of investigative heuristics for scientists [4], models of when a conclusion from observations is warranted [69], and accounts of causation that make use of intervention and statistics rather than logical implication [96].

Reduction of science to first principles. The other tenet of logical empiricism often unwittingly inherited by debates in science of security regards the unity of science or the reduction of science to single first principles. There are two senses of unity here that are not often properly distinguished: methodological unity and unity of content by reduction to a single set of models. A unity of methods would mean that, although individual sciences have distinctive approaches, there is some unifying rational observation and evaluation of evidence among all sciences. This view was deemphasized within logical empiricism. With confusing terminology, modern arguments often return to this idea under mosaic unity or pluralism: the sciences are about widely different subjects, but there are important shared social and methodological outlooks that unify science as an enterprise.

The traditional idea of reductionism is that the set of laws of one science can be logically reduced to that of another [66]. This notion requires the conception of laws as logical rules of deduction. As famously critiqued by Cartwright, the laws of physics are not true explanations of the world, but rather of the models we build of the world [8]. If laws are about models, and models can be diagrams or small-scale physical replicas, it is unclear how reduction could be defined. Bickle [5] defines reductionism (in neuroscience) as when a lower-level mechanism contains all the explanatory power necessary to intervene on a higher-level mechanism. Merits of Bickle's view aside, he has disposed of all logical-empiricist ideas of laws, deduction, and verification and uses the modern concepts of mechanistic explanation and intervention.

Reductionism is dangerous because it tends to blind us from using the appropriate tool for the job. If everything reduces to physics, then we just need a physics-hammer, and everything looks like a nail. But we shall need a more diversified toolbox in a field

[1] There is a compelling argument that 1880–1930 Vienna produced the most vital intellectual movements in science, art, and philosophy that dominated the 20th century [88].

[2] Logical empiricism is closely related to logical positivism and neopositivism; we will not distinguish these at our level of analysis [15, 89].

[3] Popper published the idea in German in 1935 though the popular English translation appeared in 1959. Carnap's work on verification, though done in 1956, is done in knowledge of and contrary to Popper. Earlier verificationists, stricter than Carnap and against whom Popper reacted, include Wittgenstein as early as 1929 [15].

such as cybersecurity. Social sciences play an equally important role as technical sciences [2]. The modern terms in philosophy of science are *integrative pluralism* [62] or *mosaic unity* [14]. The core of these terms is that fields cooperate on adding constraints to coherent explanations according to their particular tools and expertise. Such interfield explanations are what is valuable, not reductions [17]. We explore challenges due to reductionism and its alternatives further in Section 4.3 and Section 4.4.

Science as a process. A common pitfall treats the terms 'scientific' and 'correct' as synonyms. Science is a process; it yields answers. Answers are correct or not based on facts of the world. However, one calls a process 'correct' if it follows an agreed, human-defined form. The question about a process should be whether it is satisfactory in efficiently producing adequate answers. We should not assume answers are reducible to one 'correct' answer; many answers may adequately satisfy a purpose [83]. Conflating 'scientific' with 'correct' and 'correct answer' with 'adequate' results from logical-empiricist assumptions in the common complaints.

Removing these faulty assumptions is not a panacea. A sound scientific process may produce unsatisfactory results if the subject matter is difficult to study for undiagnosed reasons. One may construct a model of a system or phenomenon using scientifically rigorous methods. The model constructed will have certain properties that are considered correct if it captures sufficiently adequately the properties of the system or phenomenon that are required to address the questions that the model is intended to explore. Ultimately, the goal of this paper is to refocus the question from 'is this process scientific' to 'why is this scientific process producing unsatisfactory results'.

As we shall see in Section 3, logical empiricist threads pervade existing discussions of science of security. With this basic history and the risk to uncritically adopt these ideas both established, we shall continue to tie in the critical reflection from modern philosophy of science.

3 EXISTING STATEMENTS OF SCIENCE AND SECURITY

Many organizations have proposed problem statements and solutions regarding the state of information security research. Since 2008, these statements are frequently phrased using the language of a science of security. The motivation and goals are complex, but one important consideration is policy makers asking for intelligible explanations that can inform their decisions. In this section, we survey first the problem and then proposed solutions.

There seems to be broad agreement that there is a problem with the state of information security. That sense predates the arguments that science is the answer; education and standardization efforts predate the focus on science. More accurately, developing a science of security is part of a multi-pronged approach by the US government, later picked up by others, to respond to threats to IT infrastructure. As early as 2001, the National Science Foundation (NSF) funded both student scholarships and academic capacity building (e.g., designing courses) to universities designated by the National Security Agency (NSA) as a Center of Academic Excellence (CAE) in Information Assurance Education [68]. The NSA CAE program began in 1998. The National Institute of Standards and Technology (NIST), and its predecessor the National Bureau of Standards, have been involved in information security standards for decades. The IETF and IEEE are at least as prominent as NIST. The study of how security standards require different features than usual information technology standards has only just begun [53]. However, around 2008, there seems to have been a shift emanating from the US Department of Defense (DoD) that brought the question of science to bear on information security problems.[4]

The DoD expresses the motivation for its scientific shift in its tasking to MITRE, quoted by MITRE's JASON office in its final report. The reason for the timing is less clear, but the concern that precipitates the question of science and security is clear. This concern is worth quoting at length:

> "The Department of Defense, the Intelligence Community, and the planet have become critically dependent on the Internet for services, transactions, social interactions, communications, medical treatments, warfare; virtually everything. Cybersecurity is now critical to our survival but as a field of research does not have a firm scientific basis. Clearly, there is a scientific basis for the infrastructure of the internet such as computation, communications, and integrated circuits but its security attributes are often vague or un-measurable. ... There are concerns that future damage could be catastrophic to major components of our core infrastructures such as power and water distribution, finance and banking, even the ability to deploy forces to defend the country.
>
> Our current security approaches have had limited success and have become an arms race with our adversaries. In order to achieve security breakthroughs we need a more fundamental understanding of the science of cyber-security. However, we do not even have the fundamental concepts, principles, mathematical constructs, or tools to reliably predict or even measure cyber-security. It is currently difficult to determine the qualitative impact of changing the cyberinfrastructure (more secure now or less secure?) much less quantify the improvement on some specific scale. We do not have the tools to do experiments that can produce results that could be compared to theory, models, or simulations. Request the JASONs consider whether cyber-security can become or should be a science. If so, identify what is needed to create a science of cyber-security and recommend specific ways in which scientific methods can be applied to cyber-security. If not, what can we learn from the practice of science that would enable us to improve the security of our cyber infrastructure and assure the integrity of information that resides in the information technology infrastructure?" [63, p. 9-10]

Note three key aspects of the problem statement. First, information security is of critical societal importance. Secondly, a desire

[4] The first use of 'science of cybersecurity' or 'science of information security' is elusive. Google Scholar searches for these terms (with quotes) on Mar 1, 2017, restricted to 1990-2007, yield exactly one plausible result: a 2004 work on critical infrastructure policy [95]. Perhaps [95] borrowed from the 2004 U.S. House Subcommittee on Cybersecurity, Science, and Research & Development, which she cites. However, besides in the subcommittee title, its report does not mention 'science of cybersecurity' or security science.

to predict and measure security via "concepts, principles, mathematical constructs, or tools." Third, the purpose of this prediction is to prevent future catastrophic infrastructure damage. Science is positioned as a possible answer, but is not presumed. The real question is not whether security is a science, but "what can we learn from the practice of science that would enable us to improve the security of our cyber infrastructure."

Much government-centric or government-funded science of security work seems to accept this problem statement, including the Air Force MURI project. There is one recent voice with which to compare. The inaugural event for the working conference Art into Science: A Conference for Defense (ACoD) in early 2017 held the goal:

> "*Push the art to a science: Creating a professional discipline.* To mature our practice, we need to be able to share our methodologies in a systematic and consistent way. We have many professionals in the security industry, but do not have a professional discipline. We'd like to philosophically discuss concepts in security, push them forward, and model them" (emphasis original) [26].

Interpreting slightly, we can see this goal as a problem statement. Security practitioners cannot share methods satisfactorily. Science is taken as a way to systematize knowledge at the appropriate level of generality that it can both be shared and remain useful. Sharing generalized knowledge would support the prediction and measurement of security, as identified in the DoD statement. The two statements do not disagree, but the different focus may lead to different kinds of solutions. However, the ACoD community is too new to evaluate the results of this different perspective.

We now switch to six statements of the current status of the science of security, positioned as a method to solve the problems as laid out by the DoD. First, we examine the direct response to DoD's problem statement in [63]. We then turn to a powerful DoD agency—the NSA—in its funding priorities. Third, we examine a speech by Dan Geer, head of the Central Intelligence Agency's public venture capital firm In-Q-Tel. Fourth, we inspect the mission statement of the UK Research Institute in Science of Cyber Security (RISCS). Fifth, we study the 2016 cybersecurity strategy laid out by the President of the United States. Finally, we consider a systematization of academic knowledge recently published by Herley and van Oorschot [43].

MITRE—JASON. Although the DoD does not presuppose science as the answer to their challenge statement around cybersecurity (above), the problem statement does presuppose a conception of science. This received conception directly impacts the answers possible. For example, the report concludes "There are no intrinsic 'laws of nature' for cyber-security as there are, for example, in physics, chemistry or biology" [63, p. 79]. As we shall see in Section 4, the claim that there are laws in biology is highly contested, and the notion of unqualified, universal laws anywhere has been challenged with general success.

The implicit goal of science as putting forward unifying theories perhaps leads to the recommendation that "the most important attributes would be the construction of a common language and a set of basic concepts about which the security community can develop a shared understanding" [63, p. 3, cf. p. 15]. MITRE is responsible for several language ontologies, including Malware Attribute Enumeration and Characterization (MAEC) at the time of the JASON report.

JASON does not provide a concise statement of what science means within the report, or what the report is hoping to provide. The report searches for "guidance" from other sciences, namely economics, meteorology, medicine, astronomy, agriculture, model checking, and immunology. Thus we presume the authors judge all these fields to be sciences worthy of emulation. The report notes that sciences need not conduct experiments and may be observational. The "crucial feature" is that data be "generalizable" [63, p. 34]. Unfortunately, the report is silent on how to achieve this. The closest JASON gets to a formulation of what science of security would contain is to say "it is not simple to define what the 'security' in cyber-security means" and to call for precise definitions [63, p. 22]. One gets the sense that the authors explained what existing sciences may contribute to security, rather than how scientific methodology could be adapted to cybersecurity as an independent field.

NSA. The NSA uses a definition of 'Security Science' to guide research funding considerations. Although this was posted rather obscurely to an online community forum of security researchers, the NSA operates the forum, and the description is by the technical director emeritus:

> "Security Science – is taken to mean a body of knowledge containing laws, axioms and provable theories relating to some aspect of system security. Security science should give us an understanding of the limits of what is possible in some security domain, by providing objective and qualitative or quantifiable descriptions of security properties and behaviors. The notions embodied in security science should have broad applicability – transcending specific systems, attacks, and defensive mechanisms. The individual elements contained within security science should contribute to a general framework that supports the principled design of systems that are trustworthy, they do what people expect it to do – and not something else – despite environmental disruption, human user, and operator errors, and attacks by hostile parties. Trustworthy system design may include contributions from a diverse set of disciplines including computer science, systems science, behavioral science, economics, biology, physics, and others" [61].

This definition of science of security seems to be what the Symposium on the Science of Security (HotSoS) CFP has in mind when it refers to building "a foundational science of security" [48]. The CFP does not otherwise define science of security, but the conference is funded largely by NSA. The definition has also influenced academic work, such as that summarized in a special issue of S&P Magazine in 2011 [25].

The NSA definition defines the characteristics of an answer, not a process. The science is a "body of knowledge," it provides "objective...descriptions," and should support "principled design of systems that are trustworthy." Such goals describe the outputs of a process, not how to conduct the process so as to achieve these results. Advice would be more actionable for practitioners if it guided how to act in order to bring about an end goal instead. This

observation does not make the NSA statement wrong, just less useful than it may first appear.

Security academics informed the NSA definition. The NSA collected its influential thinkers for a special issue in "The Next Wave" on a "blueprint for a science of cybersecurity." The magazine highlights emerging trends salient to the NSA. The issue entrenches a logical empiricist position on a science of security, especially Schneider's article [60]. Like the JASON report, some articles mention biology or engineering, but nothing systematic and with no true departure from the inherited logical empiricist world view.

Dan Geer. Related to HotSoS, the NSA also runs an annual 'science of security' award for best paper. One of the distinguished experts that review the nominations, Dan Geer, provides the following insight into the reviewers' decision making:

> "Amongst the reviewers our views of what constitutes a, or the, Science of Security vary rather a lot. Some of us would prioritize purpose... Some of us view aspects of methodology as paramount, especially reproducibility and the clarity of communication on which it depends. Some of us are ever on the lookout for what a physicist would call a unifying field theory. Some of us insist on the classic process of hypothesis generation followed by designed experiments" [33].

The disagreement highlighted by this statement is that cybersecurity experts may view the area with which they are familiar as the area that makes security into science. This bias is natural, any expert has likely pursued what they view as the most important topics. However, this strikes us as strange: why would any of Geer's listed possible priorities take primacy? All contribute to a wider understanding of the world and its mechanisms via different methods. There are unifying aspects, and important differences, between different biological sciences, as one example. However, reducing the decision to one or another feature, as [33] indicates the reviewers of the best 'science of security' paper are disposed to do, collapses away the nuance necessary for a constructive perspective on a science of cybersecurity.

RISCS. For a perspective outside North America, we use the Research Institute in Science of Cyber Security (RISCS), established in 2012 and funded by the UK EPSRC (Engineering and Physical Sciences Research Council), GCHQ (Government Communications Headquarters), and BIS (Department for Business, Innovation, and Skills). Further, RISCS is cited by The Royal Society in their strategic plan for cybersecurity research in the UK generally [78]. The statement of purpose summarizes the Institute's mission:

> "RISCS is focused on giving organisations more evidence, to allow them to make better decisions, aiding to the development of cybersecurity as a science. [RISCS] collects evidence about what degree of risk mitigation can be achieved through a particular method – not just the costs of its introduction, but ongoing costs such as the impact on productivity – so that the total cost of ownership can be balanced against the risk mitigation that's been achieved. [RISCS]'s main goal is to move security from common, established practice to an evidence base, the same way it happened in medicine" [90].

The emphasis on science is much more pragmatic in this British context. The primary goal of the Institute is to provide evidence for improved decision making; this in itself is taken to advance a science of cybersecurity. This approach neatly sidesteps many of the questions about what makes a science scientific that bog down the previous discussions. RISCS issues annual reports about its work, and we may consider this to be leading by example, but it is not self-reflective about how the work advances a science of cybersecurity [76]. It is not enough to say we need to use evidence, like in medicine. That statement is true, and we agree that the work done by RISCS certainly is scientific and does advance a science of cybersecurity. Similarly, the work presented at HotSoS and otherwise supported by US DoD has a positive impact on the field. We would like to extract the reasons why. For this explanatory task, the pragmatic statement of RISCS is not enough.

White House. We return to North America and the highest levels of government. The White House has developed a 'cybersecurity research and development (R&D) strategic plan' [80, p. 2]. To a large extent, the plan takes the definition of science for granted. The plan is much more about what types of work the agencies should prioritize funding. However, these priorities explicitly include reinforcing a scientific foundation and a research infrastructure that are deemed to be lacking.

> "Cybersecurity...needs sound mathematical and scientific foundations with clear objectives, comprehensive theories (e.g., of defense, systems, and adversaries), principled design methodologies, models of complex and dynamic systems at multiple scales, and metrics for evaluating success or failure. ...[Currently,] most techniques are domain- and context-specific, often not validated as mathematically and empirically sound, and rarely take into account efficacy and efficiency. Thus, the state of the practice consists of heuristic techniques, informal principles and models of presumed adversary behavior, and process-oriented metrics" [80, p. 30].
>
> "Research Infrastructure: Sound science in cybersecurity research must have a basis in controlled and well-executed experiments with operational relevance and realism. That requires tools and test environments that provide access to datasets at the right scale and fidelity, ensure integrity of the experimental process, and support a broad range of interactions, analysis, and validation methods" [80, p. 13].

The strategic plan emphasizes certain aspects of security to focus on, for example defense should focus on deter, detect, protect, and adapt. This is unobjectionable as far as practical direction goes. However, we take these to be the subject-areas about which to do science, but not related to any definition of science. The scientific foundations and research infrastructure are cross-cutting issues of methodology to be applied to the subject matter of priority. In this way, the strategic plan plausibly separates the desire for answers from the method by which reliable answers are derived. At the same time, the prescriptions will come to seem overly rigid in our analysis. Why are domain-specific techniques not scientific? The statement "most techniques are domain-[specific]" above seems to imply that this state of affairs is unacceptable compared to "comprehensive theories," but this argument is unclear. Does the fact that specially-designed radio telescopes for finding pulsars in the centers of distant galaxies cannot be used to analyze toxicity in marsh ecosystems make finding pulsars unscientific somehow? Clearly not.

Academic Survey. Herley and van Oorschot provide a comprehensive academic perspective in a recent paper [43]. Their main thesis takes a pessimistic outlook:

> "[T]he security community is not learning from history lessons well-known in other sciences. ... What is [hard] to accept is [The security community's] apparent unawareness or inability to better leverage such lessons" [43, p. 16].
>
> "...practices on which the rest of science has reached consensus appear little used or recognized in security, and a pattern of methodological errors continues unaddressed" [43, p. 1].
>
> "...The failure to validate the mapping of models and assumptions onto environments and systems in the real world has resulted in losing the connections needed to meet [the goal of improving outcomes in the real world]" [43, p. 16].

The belief underlying this assessment seems to be that security researchers are unfamiliar with scientific methods or philosophy of science. This claim motivates an argument that the science of security initiative is essentially too vague to be useful. The solution Herley and van Oorschot advocate is to introduce the philosophical literature and draw practical lessons for security researchers. We concur with this general assessment, and the direction of the discussion and solutions. However, we believe the solution does not focus on the most relevant or helpful segments of the philosophical literature. Instead, they rely on historical framings emphasizing logical deduction and the problem of induction. This framing inherits a world view from logical empiricism, Hume, and Kant. While historically important in philosophy of science, this perspective does not provide adequate tools for solving the modern challenges of a science of security. Modern philosophy of science judges the science of security less harshly than the failings identified by Herley and van Oorschot, while providing more actionable positive advice. Here, we focus on the philosophical summary and argue why the proposed solutions in the SoK are inadequate; Section 4 provides the positive advice from modern philosophy of science.

Supposed failures to apply lessons from science are the core argument. Some are novel compared to the government and industry positions explored above; for example, "failure to seek refutation rather than confirmation" and "reliance on unfalsifiable claims" [43, p. 11,13]. Unfortunately, these observations require logical empiricist views, notably Popper and Ayer. As explained in Section 2, logical empiricism is an outdated view that has been supplanted by more recent scholarship in the philosophy of science. Relying upon logical empiricist arguments has unfortunately led the authors to draw conclusions that are often unhelpful or incorrect.

Consider the criticism of failure to seek refutation rather than confirmation. What do we refute in security? If it could be logical sentences somehow implying authorization, perhaps this is a useful criticism. However, authorization is a policy decision; in the terms of logic, it is a semantic property. We need to define the model structure, satisfaction condition, and domain before we can interpret any sentences and check semantic properties. Such definitions can be argued for or justified, but are always contextual and not something we usually talk about as refuting or confirming. Like all of security, it cannot be done according to absolutes. This logical-empiricist drive for logical absolutes confounds security just as quickly as it has confounded other sciences. A better expression of these worries is that generalizations from particulars should be properly evidenced and that reasoning from existing knowledge should be justified appropriately. We take the mechanism discovery literature as a better framework in which to discuss generalization [85]. While the philosophical aspects are not made explicit, this emphasis on evaluating evidence and warranted generalization is consistent with the arguments put forth by Shostack and Stewart [82].

Because [43] takes falsification as central, it is silent on how to draw useful generalizations in the social sciences. Since a social science perspective is needed to understand cybersecurity [2], this is a significant shortcoming. To see this, consider the statement "refuting evidence is [always definitive]" [43, p. 16]. This statement assumes the item being refuted has certain properties; namely, it must be a decidable, valid logical sentence in a sound proof system with excluded middle that is satisfied in a relevant model structure. Common, useful biology and sociology statements do not have these properties. Instead, sociologists [24] and biologists [36] tend to talk about mechanisms. We expand on such alternatives in Section 4.3.

Two other failures, "to bring theory into contact with observation" and "to make claims and assumptions explicit" [43, p. 12], are already represented by complaints from other sources about reproducibility, challenges to experimentation, and a common language. Making claims explicit is generally good advice, though [43, p. 12] wants explicitness so observations will fit into a common-language of a logical theory. Thus we wish to dispense with the common-language critique while retaining the recommendation of explicitness. Explicitness has been recommended in science of cybersecurity previously by Maxion, under the term "structure" or argument clarity [57], and by Hatleback and Spring as "transparency" [40].

Despite our criticism, many recommendations in [43, §5] are good. We agree that physics-envy and crypto-envy are counterproductive, for example. So why do we care that they rely on outdated philosophy – logical empiricism – to get there? For one, the reasons we arrive at these conclusions matter. Physics- and crypto-envy are counterproductive because there is nothing special about them to envy. From a logical-empiricist perspective of logical laws of nature, falsified by observation presenting contradiction, physics and crypto are especially well-suited. Rejecting crypto-envy would not make sense if it were actually well-suited to our definitions of science. It matters that we do not merely say 'do not worry you are not as good as physics' but instead 'physics is not as unique as you think it is'. Craver's conception of *mosaic unity* in the mechanisms literature [14] is a more useful framework to understand why crypto-envy is counterproductive. Each field participates in a multidisciplinary endeavor like security by contributing constraints on complex mechanistic explanations. Crypto is just one such field, and all fields produce constraints that are, *a priori*, of equal value to the overall explanation.

Summary. These six broad statements highlight common themes. All agree that we have a problem: cybersecurity is vitally important to society, yet not sufficiently understood to produce reliable systems. Each account on the source of that problem directly informs proposed solutions. The academic SoK uses historical philosophy of science to suggest what security needs. The other North American

statements, from the US DoD and MITRE especially, implicitly use this logical-empiricist view heavily. Dan Geer's view highlights a naïve reductionist philosophy of science closely related to logical empiricism. RISCS's pragmatic statement carries little philosophical baggage, but provides little advice on how to adequately gather and evaluate evidence. A common mistake is to confuse evaluation of the process of doing science with the evaluation of the answers the process provides.

The common thread is to look to biology, physics, or other sciences for advice. This search may be misplaced. Philosophy of science is the independent field that discusses how to execute other sciences and such issues. As we shall see in Section 4, we can disqualify all of the complaints extracted from the above that cybersecurity is not a science. Security is a science, and should look to philosophy of science to address the genuine challenges of a science of cybersecurity.

Some claim a science of security is not possible; some there is no science of security yet (see Hatleback [41] for a recent expression of this view); some just that too few people practice it. By contrast, we disassemble the argument that such a science is impossible by explaining how modern philosophy of science supports the practice.

4 PRACTICING SCIENCE OF SECURITY

We have seen a broad selection of complaints that security is not scientific enough. In this section, we contrast those complaints with alternatives based on a more comprehensive view of science according to the philosophy of science literature. The immediate aim is to clear away these unjustified complaints that security is unscientific. We do not claim, nor imply, that clearing this ground directly grants a science of cybersecurity. We must cast away these ghosts before we identify and tackle genuine challenges. Table 2 lists the five common complaints and summarizes how to defuse each with the positive advice in the philosophy of science literature.

4.1 Scientific methods

Claim: Experiments are untenable. The inability to conduct experiments, at least in many contexts, is held to be a major strike against a science of cybersecurity. The received view is of an explicit evidence hierarchy, with randomized control trials (RCTs) at the top [10, p. 135ff]. This view is rooted in medicine, influenced public policy generally, and in turn security. A critical study of proper evidence-based policy summarizes the received view succinctly: "You are told: use policies that work. And you are told: RCTs—randomized controlled trials—will show you what these are." And yet, they immediately follow with "[RCTs] do not do that for you" [10, p. ix].

Nevertheless, experiments generally, and RCTs specifically, hold a high status. High enough status that many statements from Section 3 present lack of experiments as sufficient grounds to demonstrate security is not a science. Ultimately, this high status reaches back to a conception of science as falsifying logical statements inherited from logical empiricism, for which RCTs are well-suited. We counter three common reasons for the claim experiments are untenable in security: lack of suitable control, privacy constraints, and rapid technological change.

However, first we show that untenable experiments, narrowly understood, is not the right complaint in the first place. Section 2 expanded our view of scientific explanation beyond falsifiable logical statements. Therefore, we need broader methods of evaluation than experiments, which are designed primarily to evaluate logical-empiricist-style claims.

Alternative: structured observations of the empirical world. Experiments simply are not a necessary condition for a field to be a science. We cannot induce controlled hurricanes, yet meteorology remains a science. Similarly for astrophysics—we do not induce supernova—and paleontologists, as we do not induce extinction via meteor impact. Social sciences abound that rely on case studies; an individual case is "a specific, a complex, functioning thing" [86, p. 2].

We prefer the term 'structured observation' over experiment as a necessary feature of science. Structured observations include both experiments and case studies. Robust research methods provide the structure, for example as described by [34, 86]. Structured observations are empirical, and this includes both qualitative and quantitative studies. Let us rephrase the complaint, then, as structured observations are untenable in security. Nonetheless, we will clearly demonstrate how those practicing a science of security can and already have been overcoming objections in the context of structured observations.

Overcoming: Lack of Control Groups. There are surely some situations in studying security in which randomized control trials (RCT) are not possible. Such a trial involves dividing a group of subjects such that the only statistically-meaningful difference between the two groups should be the intervention of interest. This structure permits statistical determination of the truth of the intervention's impact, granted various properties about the design hold.

Randomized control trials have come to be considered a cornerstone of evidence-based medicine, and there is prestige associated with RCTs. However, recent projects challenge this perception of RCTs as standing alone at the top of a strict evidence hierarchy. For example, without mechanistic evidence, one cannot decide the details of what RCT to design and conduct [94]. Such arguments do not cast doubt on the use of RCTs when they are plausible, but rather cast doubt on the undisputed primacy of RCTs as the *best* evidence. Various interlocking kinds of evidence are necessary for evidence-based medicine; we see no reason why security should differ in this regard. Case studies, natural experiments, model-based reasoning, and RCTs all have important, interdependent roles to play. This insight helps sharpen what is actually needed in security research to make it more like medicine, as called for by the UK's RISCS.

There are instances where RCTs have a role to play in security, particularly where security interfaces with psychology, e.g., in usable security. Examples are success in studying alternatives to passwords [6] or biometric uses of keystroke dynamics [49]. Usable security experiments do have pitfalls essentially unique to the field; for example, to make sure users have a realistic sense of risk in the lab environment [52]. And like in medicine, evidence is needed to link the experimental result to cases outside the lab.

To see a variety of approaches to structured observations in action, we now briefly review research involving passwords. Case study methods yield insights into how people select passwords and

Untenable experiments	Structured observations more broadly, not just experiments, are necessary for science. Qualitative research methods [34] such as case studies [86], and natural experiments [64], provide usable intellectual structure. Privacy and ethical concerns have been adequately addressed by the Menlo report [21]. Rapid technological change makes generalization of results a genuine challenge, but generalization tactics should help [65, 85].
Reproducibility is impossible	Reproduction comes in many forms (corroboration, statistical power, repetition, etc.) and usually several, though rarely all, work [27, 87]. The misconception is requiring all forms simultaneously, which is overkill. For a historical touch point, see [9]. Traditional scientific work sometimes covers non-replicable events, e.g., the extinction of the dinosaurs [35].
No laws of nature	'Law' interprets how scientists explain or generalize knowledge, but is too rigid even to describe physics [8]. Causal explanation as intervention is well-developed [38, 39, 96]. Philosophy of science provides access to a rich set of mechanism discovery heuristics used in other sciences [4, 14, 16] that can be productively ported to security [84]. These heuristics for designing and interpreting observations are not available with 'laws' as our goal.
No single ontology	A single language does not define a field. Within physics, the subfields communicate via trading zones in which specialized languages enable exchanges between the jargons of two subfields [28]. Trading zones apply in security as well [30]. Neuroscience provides a better metaphor for demarcating a science of security: the mosaic unity coheres from multiple subfields providing constraints on multi-level mechanistic explanations [14].
'Just' engineering	Subsuming engineering under science [83] or science under engineering [50] is not satisfying. Engineering as usually practiced depends on science [92], while at the same time science as usually practiced depends on engineering [20]. Our tentative working definition differentiates based on the goals: engineering is forward-looking, but science tries to generalize models from structured observations. By this definition, a science of cybersecurity clearly exists.

Table 2: Summary of five common complaints raised by the science of cybersecurity community and recommendations on positive actions from the philosophy of science literature to counteract these complaints.

what motivates their behavior. An example of a qualitative case study comes from Wash, who conducted in-person interviews to identify the mental models people use when thinking about security, including passwords [93]. Wash found that while everyone agreed that selecting good passwords was important, articulating why or how was much harder. Gaw et al.'s quantitative case study of 49 undergraduate students documented widespread password reuse, along with incorporating non-private attributes such as phone numbers into passwords [32]. These case studies are complemented by later observational studies carried out at a larger scale. For example, Das et al. analyzed hundreds of thousands of leaked passwords to quantify the prevalence of password reuse and other insecure activities [18]. This study corroborated earlier case studies. Motivated by this and other studies, researchers have proposed new mechanisms to enable better password selection, which can then be evaluated empirically. For example, Ur et al. ran an experiment in which users selected passwords with the aid of 14 deployed strength meters [91]. While the most stringent meters did elicit stronger passwords, they also required longer interactions and stirred greater resentment among users. Egelman et al. proposed a strength meter, then evaluated it using both a laboratory experiment and field study conducted over a much longer period for selecting a lower value password [23]. Interestingly, while in the experiment users selected better passwords using the meter, in the field experiment on low-value passwords, the meters had no discernible effect. Finally, governments are now basing recommendations for password selection and use informed by the academic literature [67].

What lessons can we draw from this brief survey through the passwords literature? First, that many methodological approaches can be used. Second, that structural observations can improve our understanding of a problem and produce better technology.

Overcoming: Ethical Constraints. A further concern is that experiments are impracticable in security for privacy and ethical reasons, rather than simply being impossible to design properly as the foregoing argument held. The ethical considerations of security studies have been traced in detail by the Menlo Report [21]. In the biological and social sciences, the Belmont Report established the ethical guidelines by which experiment design should be held, in the US and UK some of these guidelines have been put into law. Universities enforce these policies via review boards that oversee experiment design before the experiment can be run. The Menlo Report updates the three classic considerations – respect for persons, beneficence, and justice – for an internetworked world. A fourth, respect for law and public interest, is added. Privacy plays in all four of these organizing principles.

Ethical restrictions are a basic part of research in many scientific fields. Neuroscientists cannot open up the brains of human test subjects and apply electrical shocks. Physicists should not wantonly release radioactive particles to test atomic phenomena. Virologists cannot test the spread of disease by releasing a pathogen in an airport and tracking its progress. All these fields make do by designing ethical observations that get to the necessary explanatory insights. We already have a thorough update to these ethical considerations

for information and communication technology in the Menlo Report [21]. Engaging with ethical review boards may slow down security researchers at first. But it has not stopped other sciences, and it should not stop security from being viewed as a science.

There is some nuance to this privacy challenge to experiments, which is that participant data being private means that experiments are not reproducible; the data cannot be shared with other researchers. Using our language developed in Section 4.2, this is only a problem for rerunning statistical tests. In the other seven senses, the data is re-collected. If all the other artifacts for experiments are available, including the code to collect and analyze data, repetition, reproduction, and so on should be possible without knowing the original participants. And even then, in many cases the original data may be anonymizable. Therefore, the cost in terms of reproducibility for the researcher to comply with ethics and privacy appears to be minimal.

Overcoming: Rapid Change. The third critique on the tenability of structured observations in security concerns the longevity or generalizability of results due to the pace of technological change. The critique runs, roughly, that although experiments are plausible, their results are not useful because the results are outdated by the time they are compiled and published.

This critique rests on a combination of selecting the phenomenon of interest and a conception of who culturally is doing science. Some phenomena are more ephemeral than others, in the sense that the phenomenon only exists for a short time. The extinction of the dinosaurs was an ephemeral event in this sense [35]. Ephemeral is not to be conflated with quick. Chemical reactions may happen fast, but the phenomenon of gunpowder exploding is a robust chemical phenomenon, for example. If we select an ephemeral phenomenon of interest in security science, we cannot be surprised that the resulting explanation has short-lived relevance. We are especially likely to do so when doing forensics, a key function within security, because to investigate the specific details of past events is almost always to investigate ephemeral mechanisms. This focus is no different from the difference between paleontology and contemporary biology, for example.

The selection of ephemeral phenomena interacts strongly with who is considered a scientist in the security landscape. Beyond universities, many government organizations and private companies have security research labs. Even beyond staff with 'research' in their job description, our notion of engineering as satisficing to create artifacts and science as crafting generalized knowledge induces some interesting perspectives. A software developer writing code is building something, therefore engineering. But this distinction no longer applies when the developer is doing a code review of a colleague's code, or trying to find the point in the code responsible for a bug. Surely, the developer is following established practice in these tasks, not inventing new modes of experiment [72]. But we do not say that a high-school student following textbook steps for a chemistry experiment is not participating in the scientific enterprise, just because the steps are known. The line blurs further for older students, following known procedures but with slightly varied chemicals in order to measure any potential differences. A developer may follow textbook steps for establishing a hypothesized source of a bug, making an intervention in the system, and taking measurements to establish evidence for whether or not the intervention confirmed the hypothesis. But the code is new, so the answer is not known a priori.

We consider developers, or more likely for security malware reverse engineers investigating the functionality of unknown code based on hypotheses, operating in this mode to be participating in science. Therefore, the longevity-of-results picture changes. Experiments may have a short window of application, but cybersecurity experiments may also be quite quick to execute and apply. An experiment might be automated malware reverse engineering that runs the malware in a sandbox, extracts suspicious domain names and IP addresses from connection attempts, and publishes those network elements to a blacklist. The time frame between the beginning of the experiment and operational impact on defender networks may be 15 minutes. Just because it takes the US FDA 15 years to approve a drug does not mean anything scientific takes years. The network elements may have a short lifetime of usefulness, but it is at least proportional to the duration of the experiment.

The critique that experiments in security are untenable because of rapid technological change takes an unacceptable combination of options – that these miniature experiments are not science, but that science must cope with highly ephemeral phenomena without such experiments. There is no reason to conceptualize science in this contradictory way. In either other conception of science's relation to automation and the expectations we have for its results, the challenge that rapid technological change makes it impossible simply evaporates. There may be methodological or generalization-based challenges to a science of cybersecurity so conceived, but we can confront them directly once recognized.

Structured observations via mathematical modeling. Another critically important purpose of structured observation is to support the construction of mathematical models of systems, perhaps incorporating representations of policies. The use of mathematical modeling in supporting the practice of science and engineering is well established. The process by which mathematical models are constructed can be summarized conveniently by the diagram given in Figure 1. Observations of the world are performed; by a process of induction, the construction of a model is commenced; using deductive reasoning, the properties of the model are derived; those properties are interpreted in the observed world; and further (structured) observation of the world is used to assess the accuracy or utility of the model as constructed so far. This process of observation, construction, deduction, and interpretation is iterated until the modeler is satisfied that the constructed model is fit for purpose. This process by which mathematical models are constructed makes use of both inductive reasoning, in which we draw conclusions about the empirical world using observations and inferences from those observations, and deductive reasoning about the mathematical model itself. The process seeks to constrain each type of reasoning by reference to the other.

Both during the construction of a model and during the use of a completed model in an engineering process, two kinds of reasoning about the properties of models are used. First, models can be explored intensionally, in the style of experimental mathematics; that is, the space of evolutions of the model is explored systematically using simulation methods, such as Monte Carlo [12, 46],

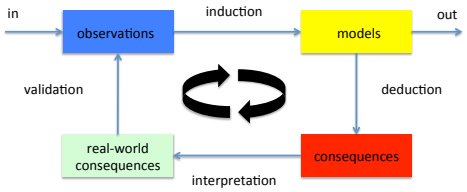

Figure 1: The mathematical modeling cycle

and absolute and expected properties can be observed and deduced. For example, [11] explores the consequences of different levels of resourcing for the effectiveness of access control to buildings and consequent information security breaches. Second, properties of the model can be explored extensionally; that is, the full range of logical and mathematical tools can be used to reason about the properties of the model considered as an object of mathematical study. This latter form of reasoning can be illustrated in the context of program verification by the set-up of Separation Logic [45, 75, 77]. There it can be seen that the natural mathematical model of computer memory — derived by direct observation of the architecture of computers — can also be understood as a model of a certain kind of mathematical logic that is appropriate for reasoning about resources (e.g., [13, 31, 70, 75]), with the consequence that full range of tools from mathematical logic can be brought to bear on understanding the behavior of programs.

We would argue that the mathematical modeling method applies in the science and engineering of security just as it applies, for example, in civil, electrical, and mechanical engineering and their supporting sciences.

4.2 Evaluating Results

Claim: Reproducibility is impossible. This complaint requires considerable unpacking. There are some intuitive challenges; for example, if reproduction of phenomena under controlled conditions is an absolute requirement, then astrophysicists and paleontologists are not scientists. This complaint inherits a flavor from logical empiricism; reproducible experiments are necessary to repeatedly test and increase confidence in falsifiable (but not yet falsified) statements of universal law. Fragile or contextual conclusions in biology—conclusions that were not readily reproducible—historically led to serious claims that biology was not a science [8].

Complaints of irreproducibility, at heart, strike out at the genuine observation that conclusions in cybersecurity research are often fragile or contextual. Philosophy of biology countered similar logical-empiricist attacks by creating a more nuanced idea of evaluating explanations and results. Philosophy of security shall leverage this work about biology to do the same. Reproducibility is a complex term in itself; we find no fewer than eight different senses of the term that discuss different aspects of evaluating evidence from structured observations. One way to view these different senses is as selecting a perspective at which phenomena are not too fragile.

Alternative: Evaluation takes many forms. Although the distinction between replication and repetition is not new [9], recent work provides actionable advice to scientists. We focus on the family of five terms suggested by [27], plus the notion of statistical reproducibility [87]. We discuss three distinct senses of statistical reproducibility, for a total of eight distinct methods to support the robustness of evidence. When one complains that cybersecurity lacks reproducibility, usually what is meant is that one or two of these eight senses is impossible. All sciences similarly struggle to achieve reproducibility in all these senses at once. Thus, cybersecurity is no worse off than other sciences.

Feitelson suggests five distinct terms for a computer science discussion of evaluating results [27]:

> Repetition – to rerun exactly, using original artifacts
> Replication – to rerun exactly, but recreate artifacts
> Variation – repetition or replication, but with a measured intentional modification of a parameter
> Reproduction – to recreate the spirit, in a similar setting with similar but recreated artifacts
> Corroboration – to aim to obtain the same or consistent result as another study via different means

Each of these strategies has its uses, one is not strictly preferred over the others. Throughout this subsection, we use an extended example of evaluation of intrusion detection systems (IDS). One may question whether evaluating an IDS rule is scientific, in the sense desired. IDS rules may be very specific and not widely generalizable, but the same could be said for determining whether a particular enzyme in the blood of some rare Amazonian fish actually selectively binds to some specific parasite.

Repetition. Even in the restricted context of evaluating a single IDS rule, these strategies all have a sensible interpretation. If we have a recorded stream of traffic, if we play the stream back through the same IDS with the same network configuration, the same rules should fire. A failure of repetition would be indicative of race conditions or performance bottlenecks, or perhaps an architecture with multiple distinct IDS hardware and a network switch that randomly assigned packets to IDS interfaces, meaning any rule that required correlation between packets would not reliably fire across repetitions. It is because of repetition experiments such as this that Bro, Snort, or Surricata work on flows, not packets. When load-balancing any of these IDS tools, traffic is balanced based on flows, so that all packets in one conversation go to the same thread. Flow-based balancing is needed because a network IDS works to reassemble application-level, or at least transport-layer, information [79]. Any network architecture that randomized packets among IDS instances, as opposed to maintaining flow-based balancing, would lead to un-repeatable observations because application-level reconstruction would fail to be the same, and the IDS tests would be applied to different information.

Replication. Replication might mean to use a different IDS software and re-write the rule in its language to get identical detection. Replication might also mean using the same IDS architecture and rules on traffic that is artificially generated to have the same malicious features as a prior experiment; or perhaps it is simply to recreate the same network architecture at a different physical location and replicate that the IDS works on a pre-recorded traffic

stream to provide evidence the newly setup sensor architecture has been installed correctly.

Let us pause here; we have claimed that creating a new IDS rule and testing it on a pre-recorded traffic stream is an experiment. Are we abusing the use of experiment here, in a way that is not consistent with other sciences? We think not. Certainly, the objects of the experiment are artifacts, and software artifacts at that. But if cybersecurity is a science of anything, certainly it is of software (and how people interact with it). Therefore, the IDS analyst making a rule has an untested specimen (the IDS rule), a hypothesis about how it should behave (what she designed it to do, in this case), and establishes a controlled environment in which to test the properties of the specimen of interest. This matches all the usual hallmarks of an experiment.

From repetition and replication, variation is straightforward and we will not discuss it in detail.

Reproduction and Corroboration. Reproduction of IDS alert efficacy is something practitioners measure often; with the same rule in a different network architecture, they evaluate the outcomes. Corroboration is similarly useful. Perhaps against different traffic patterns, do different IDS systems see similar but distinct alerts about exploitation of Heartbleed, for example, that allow us to corroborate wider claims about internet-wide abuse of the vulnerability.

Statistical properties. Stodden distinguishes several failures of statistical reproducibility [87]. Lack of statistical reproducibility may result from poor design of observations, namely issues with low statistical power or sampling bias. Even if the sample design is appropriate, the statistical tests applied may be inappropriate. For example, a test may be applied despite requiring assumptions that the situation does not meet. Finally, results may be generalized beyond the bounds of what statistics justifies. A rough translation of these three problems is with design, analysis, and interpretation of observations.

If an observation is poorly designed, it will not be able to be reproduced. Poorly designed means either there is a consistent, unknown confounding factor in the design (sample bias) or that the number of elements is not large enough to produce results independent of natural random variation. Sampling bias happens quite naturally across various organizations that might deploy the same IDS rule – different adversaries attack banks than attack militaries than attack universities. In security, this reduces to a challenge of comparing like with like. In medicine, it is understood that patients from different age groups or income brackets have different health outcomes. But these differences are measured, and then can be controlled for when observing the impact of a treatment on a varied population. The sampling bias is controlled by knowing the shape of bias with sufficient precision. Security suffers from an additional layer of evidence-gathering bias. Organizations may not report or disclose vulnerabilities, for example, due to ignorance, fear of financial risk, legal obligations, or to improve the reputation of their brand [82, p. 52ff]. Such social and cultural biases apply to many types of security evidence. Adequately understanding these biases, and how to mitigate them, remains an area for further work.

The statistical analysis and tests performed after data are collected can also impact whether a compatible result can be obtained in any of the five strategies for confirming results. One common problem is for researchers to selectively report results, and even tests performed. If a researcher runs "many tests until one of them yields a p-value of less than 0.05 and then report[s] this result as if the other tests had never been run" then the result is actually a statistical outlier, even though it is being reported as a reliable result [87, p. 6]. The researcher essentially misrepresents the robustness of the result, which of course impacts confirmation attempts. A science of cybersecurity must guard against such misapplications of statistics just like any other science.

The final area of statistical reproducibility we discuss relates to the data collection and generation process. Problems with data generation lead to a lack of generalizability of the results. For example, studies commonly report a linear regression establishing a relationship between two measurements, such as deploying an IDS rule and the number of intrusions. Strictly speaking, unless the data generation is very strictly constrained, one may not safely use that relationship anywhere outside the data collected – the result may not be generalized [87, p. 6]. Generalizability is indeed a problem in cybersecurity. However, to the extent that this challenge is due to failure to meet known statistical constraints on data collection, cybersecurity is not distinct from other sciences.

Forensics and Reproducibility. One intuitive objection about reproducibility stems from the idea that security is forensic, and so necessarily considers single events. By definition, if some security event only happened once, it cannot be repeated. We treat forensics as importantly different from science (see Section 4.5); however, there is not a sharp distinction. Establishing evidence for a particular series of events at a particular time in the past is forensics. Sometimes sciences require such evidence as part of knowledge creation, generalization, and application to new scenarios. It seems implausible to claim that astrophysics, paleontology, and macroeconomics are unscientific. Yet they are largely historical, in a way very similar to digital forensics.

A science of security can integrate forensics in an analogous way to how biology integrates paleontology and physics integrates astrophysics. Paleontologists build evidence for the mechanism of what killed the dinosaurs, for example. Glennan refers to these one-time events of interest as being driven by "ephemeral" mechanisms, as opposed to say chemistry where mechanisms have a more robust, repetitious nature [35]. The reason for unifying ephemeral mechanisms, as studied in paleontology, with other sciences is because mechanism discovery strategies and mechanistic explanation provide a coherent account of scientific activity. Bringing paleontology into that fold brings the philosophical account of scientific explanation via mechanisms in line with the obvious similarities between astrophysics and paleontology on one hand and physics and biology on the other. Biology is not unscientific because it contains a sub-field focused on single events in the past – paleontology; similarly, a science of cybersecurity is not totally scuttled simply because it contains a sub-field that focuses on forensics.

Summary. The thrust of this argument is that observations in cybersecurity appear, to a large extent, to be amenable to the various senses of reproduction. We consider these senses as evidence evaluation terms. However, we have not attempted to refute the fragility of conclusions in cybersecurity, as this fragility appears genuine.

We believe the question can be more productively answered by selecting the correct level of analysis than naïvely insisting against reproducibility.

Questions of reproduction skip the important issue of what phenomenon is to be reproduced. Which phenomenon is of interest will impact which evidence evaluation strategies (repetition, statistical tests, etc.) are most valuable. We as scientists are generally interested in discovering the mechanism responsible for a phenomenon in order to better explain the phenomenon. Defining the phenomenon differently will change the mechanism of interest. For example, most models of computer network attacks have a step for exploitation of the target. Attacks are a high-level phenomenon and exploitation is one activity within the mechanism. At a lower level, there are various mechanisms by which exploitation could occur, for example drive-by downloads or social engineering to run untrusted code [84].

4.3 The nature of scientific inquiry

Claim: No laws of nature. The critique that there is no science of cybersecurity until there are laws of security, or mathematical rules which permit deductions from observations to consequences, comes presumably from the received view that this is how physics works. The importance of laws in the received view is to provide predictive power. However, we go too far if we claim we need laws in order to make reliable predictions. Many sciences make reliable predictions without laws; in fact, many philosophers argue physics does not have laws in the sense commonly understood. Yet, many predictions of physics have a reliability that surely a security researcher envies. This section introduces philosophical perspectives on how to create and evaluate predictive models.

Some science of cybersecurity work has noted that not all sciences have laws. The Army Research Lab is more thorough, pragmatically defining a science of security by specifying its subject matter, and taking models generally as central [51]. However, no prior work goes nearly far enough in breaking the preconception about laws nor providing alternatives. We shall do both. First, biology has faced and refuted this conception of laws being a necessary criterion for science in detail; we provide a brief summary. Second, we adapt to security the modern conception of explanation that has supplanted that of laws – mechanistic explanation.

First, let us summarize of what a laws-based explanation is intended to mean. The DoD takes laws as "the basis of scientific inquiry" in its request to MITRE [63, p. 4]. Hempel provides a concise definition that provides a historical basis: a law is "a statement of universal conditional form which is capable of being confirmed or disconfirmed by suitable empirical findings" and a law, as opposed to a hypothesis, refers to such a statement that "is actually well confirmed by the relevant evidence available" [42, p. 35]. In this 1942 definition we see all the famous features; for example, we see Popper's falsifiability criterion in different words ("capable of being disconfirmed"). We pull hard on this logical-empiricist, laws-based conception of explanation with a detailed unpacking of the meaning of the term followed by critiques by the modern philosophers of science Mitchell, Cartwright, Bogen and Woodward, and Woodward following Pearl.

To understand 'law' in this sense one must focus on the technical-philosophical meaning of three terms: universal, conditional, and capable of being confirmed. Universal means that the statement of law applies to all things, at all times, without exception or additional precondition. Conditional means an if-then statement in classical first-order logic. For a statement to be capable of being confirmed or refuted, it needs to have semantic content, or meaning, and be about the universe in which we live. The challenge here is captured by the aphorism 'all models are wrong, some are useful.' While models and statements may be semantic, they are also necessarily simplifications of our universe and there are always conditions in which that simplification is wrong. But how can a simplification be confirmed or refuted – it may be useless, but that is a far different thing than absolutely refuted. Many logical empiricists side-stepped such questions by saying that laws were true universally, so they must be independent of human artifice or language. This history makes it particularly strange to ask whether there are 'laws' of man-made system security.

Mitchell, a philosopher of science, has deconstructed a laws-based, unity of science. For example, she argues that "nature is complex and so, too, should be our representations of it" and that complexity and the resulting pluralism of models "is not an embarrassment of an immature science, but the mark of a science of complexity" [62, p. 115]. She advocates that the relevant aspect is not how theories are defined, but used. Thus, any theory that functions as an effective generalization might pragmatically be called a 'law' in physics, or a 'model' in biology. What we should care about is effective methods for generating and scoping generalizations so we know where and to what they apply.

Cartwright identifies various problems with laws-based conceptions. These include that usual laws-based conceptions cannot make sense of causal statements, and that laws explain the behavior of mathematical models of the world, rather than the world directly [8]. She calls this a simulacrum account of explanation.

Further deconstructing the logical positivist influence on philosophy of science, Bogen and Woodward [7] identify the mediating influence of data and observation on our ability to make theories. We care about the phenomena in the world, and our theories apply to phenomena. However, we observe data, and data collection is mediated by tools we build and their norms of use. The canonical example is the melting point of lead. We do not observe lead melting at 327°C. We observe many thermometer readings, with both known and unknown equipment errors and limitations, from which we statistically derive the value 327°C with some acceptably small margin of error and assign it to the phenomenon of lead melting. The argument goes on, with some nuance, that there is no law or even single theory that explains lead melting at this temperature, much less any of the individually-observed thermometer readings. The whole apparatus of experiment, including the engineering of the tools, the statistics, the metallurgical purity of the sample, and so on, are all necessary to explain the phenomenon. This perspective accords with Mitchell's proposition that complicated theories are a sign of maturity, not immaturity; as well as Cartwright's simulacrum account of explanation.

Woodward provides an alternative account of causal explanation to supplant a laws-based account. This account is known as an interventionist account because, roughly, it is based on the idea

that the only way to determine causation is by an intervention that could, in practice or in a thought experiment, make a change [96]. Woodward relies on Pearl's statistical account of causality [73], which has been updated since Woodward's treatment but with some modification is a sound statistical approach and language for discussing causation [19]. Causal explanation is not about laws of nature, but about building an explanation of the organization of elements of a phenomenon in such a way that one may intervene reliably on the outcome by changing the elements. Again, like Mitchell, there is a theme of what makes an adequate generalization of a system.

Alternative: Mechanistic explanations of phenomena. The most convincing account within the philosophy of science community is that of a mechanism [36]. The current consensus definition is that "a mechanism for a phenomenon consists of entities (or parts) whose activities and interactions are organized so as to be responsible for the phenomenon" [37, p. 2].[5] This mechanistic conception of the result of scientific reasoning is useful because it provides a basic structure on which we can build and borrow mechanism discovery strategies.

The literature on mechanism discovery strategies examines how scientists develop hypotheses to test, and the constraints they build into experiments and observations in order to test them. Mechanistic thinking is more helpful than a laws-based approach because it provides hints as to what to do and what to look for in order to build a useful explanation. Bechtel and Richardson base their initial strategy of decomposition and localization on Herb Simon's work. Their work is specifically positioned as a strategy for scientific discovery "well suited to problems that are relatively ill defined, problems in which neither the criteria for a correct solution nor the means for attaining it are clear" [4, p. 7]. This description is encouraging for security practitioners beset with ill-defined problems.

The other main contributor the mechanism discovery literature is Darden. She provides strategies used on slightly better-defined problems. If a general mechanism is known, but details of a particular entity or activity are hazy, a reasonably hypothesis is to attempt to resolve the hazy element more clearly ('schema instantiation'). If a mechanism is understood, but not its set-up or termination conditions, we can chain our reasoning either backwards or forwards from the known mechanism to constrain our hypotheses about what must exist either before or after the mechanism we know about [16, ch. 12].

Hatleback and Spring have begun to adapt mechanistic reasoning from these other sciences to computer science generally and security specifically. In [40], they discuss and resolve the apparent difficulty of discussing mechanisms that have been engineered, or created by humans, such as computer code. This different origin does not present any conceptual difficulty in understanding the function of the code as a mechanism. Like with any differing fields the exact tools used to examine mechanisms in computing would not be identical to biology, any more than radio telescopes for stars are useful in investigating mechanisms of frog ecology. In [84], they present incident response and intrusion analysis as a kind of mechanism discovery task. They show the heuristic of schema instantiation, from Darden [16], to be analogous to the heuristic that an incident responder uses when resolving the 'exploitation' step in the kill chain to a drive-by download, for example, and then again when the drive-by download mechanism is instantiated to clarify what particular domains and web services participated in the exploitation and could be blocked.

Science of cybersecurity stands to benefit from refusing to consider explanation as laws-based and instead focusing on scientific investigation as mechanism discovery. The question of whether there are laws of cybersecurity is fundamentally the wrong question to ask. Both philosophy of science and mathematics have better perspectives on generating intelligible explanations of phenomena than a laws-based explanation. The modern philosophy of science literature provides heuristics for mechanism discovery that should be helpful in orienting scientific investigations in security. The mathematical modeling cycle described in Figure 1 provides a complementary heuristic process.

When using mathematical models, the situation is clear. We identify not laws but rather properties of models. The properties of a model are used to assess its value and to support its refinement.

4.4 Scientific Language(s)

Claim: No science without a common language. In this section we argue against the idea that a single language or ontology of security would be a defining feature of a science of cybersecurity, although we agree that clarity of expression is necessary. Our main departure point is the insistence on unification into a single language, and advocate instead for a kind of integrative pluralism—which is what actually exists in all other major sciences anyway.

JASON identifies a "common language" as the "most important" attribute necessary for a successful science of cybersecurity [63, p. 3]. The other statements we reviewed are less direct, but there is a similar drive towards unification. Phrases that we interpret as sympathetic to unification of language and explanations include "contribute to a general framework" from the NSA's definition, "comprehensive theories" as opposed to "domain- and context-specific" ones in the White House plan, and the negative tone in which Geer relays the diversity of opinion among the best paper judges.

The implicit assumption amongst these statements is that, at least within security, a variety of disparate phenomena can be reduced to a relatively compact set of definitions and statements. In [71], the authors cite the JASON reasoning and make this desire to compact the semantic space explicit. The traditional statement of unification into a common language was 'reductionism,' as discussed in Section 2 to be logical deduction of one field's laws to those of another [66].

This reductionist idea of a common language implies an explanatory hierarchy. It seems more realistic to admit that explanations are about different topics, and each topic develops its own specialized language. A variety of fields contribute understanding to security mechanisms, from economics to electrical engineering to elliptic curves. We need translations between these fields. But that does not create a common language any more than translating between German and French creates a combined language; it just creates a translation.

Alternative: Trading Zones, Pluralism, and Mosaic Unity. A more convincing account of interdisciplinary collaboration comes from

[5] Past definitions with which to compare are [4, 44, 56].

physics. Galison, a historian of physics, borrows the anthropological term *trading zone* to describe the contact between subfields of physics, such as experimental and theoretical particle physicists [28, 29]. The analog is in merchant towns, cultural anthropologists observe people from different cultures coming together and creating just enough of a shared language such that commerce can happen. As commerce grows and the communities benefit, the trading zone becomes more robust. This occurs linguistically as well as materially. Galison's insight is that the same process happens between subfields of physics. There are members of each community who specialize as traders, go to places where the communities interface, and develop specialized languages (pidgins, creoles, etc.) that are incomplete mash-ups of each traders' home language. Theoretical physicists and experimental physicists do not, in fact, speak one common language. They speak importantly different languages, with their own jargons and evaluations of what features of explanation and evidence are most important. However, the two subfields can productively exchange ideas because there is a developed trading zone where ideas can be translated into a shared language from both directions and then re-translated back out to the respective communities.

Galison's insights on trading zones and subcultures seem to have been misunderstood by JASON. In his presentation about science in cybersecurity, Galison writes "in these choices of basic objects and their relation lie the basis of separation of subcultures and the robust power that that division offers" [30, p. 20]. Partiality of language is how Galison envisions these various subcultures of security communicating, just as in physics. The sense in which we need a common language is that we need various trading zones in which security researchers can communicate. We are not to wait for a single common language to emerge which then we can all speak unambiguously. It's not that such a goal may be slow or arduous to achieve; more importantly, it fundamentally undermines the robust power that division in specialties offers.

In biology, Mitchell has argued against reductionism for what she calls *integrative pluralism* [62]. For Mitchell, the units of science here are theories, understood as idealized models of various phenomena. There are many models, hence 'pluralism,' and models are neither totally isolated from each other nor usually reducible so that one supplants another [62, p.192]. Since each model comes with its own terms and specific usage of terms, if we could produce a common language that would be tantamount to unifying all our theories within security. Integrative pluralism, as a model of biology at least, informs us that we should not expect this unification of terminology to be possible except in localized, purpose-built contexts—that is, trading zones.

The most convincing analog for defining a field of science of cybersecurity comes from Craver's description of the *mosaic unity* of neuroscience [14]. The subfields making up neuroscience collaborate by adding constraints, based on their own individual methods and viewpoints, on mechanistic explanations. Like the stones in a mosaic, each subfield has its own unique structure and identity, but if we step back we see each stone contributes to a bigger picture.

Security has a similar arrangement of diverse fields contributing constraints on explanations. Economics constrains explanations of what users can be asked to do based on how they spend their money in situations of information asymmetry. Usable security constrains explanations of what users can be asked to do based on how they spend their attention. Craver is more helpful than Mitchell for security in that, for Craver, the explanations are mechanistic explanations, and that structure allows him to elaborate on how the subfields are interrelated and provide constraints. Different fields work on different parts of a mechanism or on a different level of a mechanism. There are a plurality of mechanisms, which are idealizations as theories are for Mitchell, and mechanisms are integrated via levels of explanation.

These three related conceptions of communication and explanation in science all go against the need for a common language for a science of cybersecurity. All three also support the importance of clarity of expression. If context is important, because there are a plurality of theories and contexts, it is vital for researchers to make their assumptions clear and use terms consistently. Pluralism is not an excuse for sloppiness of explanation. If anything, it militates for the importance of thorough, careful explanation. Perhaps the push for the importance of a common language in security is actually a push for clarity of expression. Certainly, methodology sections of published papers are vital. Maxion [57], for example, has argued strongly for structure and clarity as an aid to good science. However, as Maxion does, structure and clear expression are a separate problem that should not be conflated with the idea of creating a common language before science can commence. As we have demonstrated, sciences do not operate on a single common language, and it is misleading to pursue a common language for its own sake.

4.5 Engineering or Science?

Claim: Security is 'just' engineering. One might ask, if the goal is to produce reliable systems, why discuss science at all? Producing systems to a certain standard of usefulness is engineering. While engineering certainly leverages scientific knowledge, it also uses other kinds of knowledge [92, p. 229]. Indeed, the government call for a science of security looks similar to Anderson's description of security engineering:

> "Security engineering is about building systems to remain dependable in the face of malice, error, or mischance. As a discipline, it focuses on the tools, processes, and methods needed to design, implement, and test complete systems, and to adapt existing systems as their environment evolves" [1, p. 3].
>
> "[Security engineers] need to be able to put risks and threats in [context], make realistic assessments of what might go wrong, and give our clients good advice. That depends on a wide understanding of what has gone wrong over time with various systems; what sort of attacks have worked, what their consequences were, and how they were stopped (if it was worthwhile to do so)" [1, p. 5].

We resist both the assertion that there is no science to be done in cybersecurity as well as the pejorative connotation of placing engineering somehow below science. Sciences that study technology or human artifacts do have their own nuances and distinct concerns [58], but this is no reason to collapse security entirely under engineering any more than it would make sense to collapse

medicine entirely under engineering just because it is not pure biology.

Mathematical models are a key mechanism by which science is transmitted to engineering practice. A forensic investigation of a computer security incident considers a mathematical model of the computer security system and explores what sequence of events, as represented in that model, led to a certain observed outcome, as represented in that model. Thus the efficacy of the forensic investigation depends on the model providing a sufficiently good representation of the system and its associated security policies.

Alternative: Science, Engineering, and Forensics. Efforts have been made to draw boundaries and classify the relationship between engineering and science. In the two extreme viewpoints, Simon tries to subsume engineering under science as a science of design [83], and Koen tries to subsume science under engineering as a special case of heuristic problem solving [50]. There are also more nuanced views. Engineering as usually practiced generates understanding and depends, in part, on science [92]. At the same time science as usually practiced generates artifacts and depends, in part, on engineering [20]. Therefore, it seems unjustified to place science over engineering or vice versa. If we are going to engineer secure systems, in the sense of Anderson [1], we will need a science of cybersecurity to extract and generalize the knowledge with which we will build. Both building artifacts and extracting knowledge will face their own challenges, making it sensible to distinguish the tasks. Security should continue the traditional, close interplay between science and engineering.

For our purposes, we will make the following distinction between science, engineering, and forensics. This distinction is tentative, but will help us move the conversation forward by making sense of the difference between science of security and security engineering. Engineering is fundamentally forward-facing. It is about using models to build systems that satisfy certain goals in the future. Scientific practice produces more general, more shareable knowledge.[6] Such scientific knowledge is often expressed as models, but it also may be a catalog of well-documented cases. Forensics is about the past, or is backwards-facing. It is about reconstructing a historical explanation for some event of interest that we know occurred, but not how or why.[7]

We leave for future work a more detailed exposition of the relationship between science, engineering, and forensics. But as an example, we are sympathetic to Leonelli's conception of scientific understanding:

> "Understanding can only be qualified as 'scientific' when obtained through the skillful and consistent use of tools, instruments, methods, theories and/or models: these are the means through which researchers can effectively understand a phenomenon as well as communicate their understanding to others" [55].

[6]Dear [20] exposes a duality within science: between natural philosophy and instrumentality as twin, mutually indispensable, explanatory strategies for making nature intelligible. This duality blurs a divide between science and engineering, one reason our discussion here is pragmatic but tentative.

[7]One might say historical investigation, rather than forensic, unless the event is some sort of (security) policy violation. Glennan [35] discusses paleontology as a science that also extensively employs historical explanations; for our rough three categories, we would include paleontology as a largely forensic enterprise.

Even without detailing the precise relationship, clearly science, forensics, and engineering interact tightly. When our systems break, we conduct forensics to learn why and how. We then employ science to update our knowledge, improve our models, or document edge-cases based on this why and how. Adequate updates may include further, purpose-designed structured observations. We then employ engineering to adapt this new knowledge to build a better system, less likely to break. Therefore, we have a feedback loop from engineering to forensics and science back to engineering which contains no sharp distinction between a science of cybersecurity and security engineering. Our focus is the scientific enterprise,where science is understood as generalized-knowledge, evidence-based, explanation-generation activities.

5 CONCLUSION

In general, the thesis we forward is that security is, as practiced, a science with its own unique challenges. This statement contrasts with the surveyed views, which posit that whatever makes security hard also makes it a qualitatively different sort of enterprise than science. These detractors often accidentally over-emphasize some scientific field in conceiving science generally. Of course security is not particle physics, nor molecular biology. This conception of science is too narrow. This overly-narrow view can, in may cases, be traced back to outdated views related to Logical Empiricism.

The common complaints against a science of security are: experiments are impossible, reproducibility is impossible, there are no laws of nature, there is no common ontology of terms, and it is 'just' engineering. We have forwarded more effective perspectives on all these complaints that already accommodate security: structured observations of the empirical world, multiple methods for evaluating evidence, mechanistic explanation of phenomena, specialization necessitates scientific translation, and the interplay between science, engineering, and forensics.

We argue that cybersecurity suffers from the same sorts of challenges as other sciences. It is not qualitatively different. However, different fields of science are defined, to a large extent, by the characteristic challenges of their subject matter and how those challenges are approached. We should not confuse a search for interdisciplinary collaboration with a search for how to execute a science of security. Cybersecurity must learn from challenges common with other sciences while at the same time pushing forward with novel solutions to those challenges and approaches in fact unique to cybersecurity.

Also like other sciences, a science of security faces important social questions. We largely leave these for future work. Three candidates to explore are the gap between research and practice; who are the customers or recipients of knowledge produced by a science of security; and how the secretive nature its practice alters the science being done. Both Dykstra [22] and Metcalf and Casey [59] attempt to narrow the knowledge gap between practitioners and scientists; but the nature and social function of the gap should also be studied. Some customers are policy makers; future work would likely build on Jasanoff [47]. Perhaps some knowledge customers are practitioners, but as Vincenti [92] argues, academia also receives knowledge from practitioners. Systemic secrecy has caused different scientific practice in the case of biological weapons

development [3]; something related may happen in information security. We might ask how students with a classified PhD thesis differ from those with a publicly published thesis, for example.

We view challenges in security as challenges to building generalized, sharable knowledge. In many cases, the science of cybersecurity community has hinted at this conception of the challenge. Generalization is woven through the discussions of the difficulty of confirming observations, designing experiments, and developing a common language. Generalization is implicit in the discussion of laws because, traditionally, laws are a formal vehicle for expressing generalizations. These descriptions may accurately identify that generalization is hard in science of cybersecurity, but the diagnosis of the cause of this challenge misses the mark, as Section 4 demonstrates. This conception of generalization as the core problem of a science of cybersecurity makes particular sense with our tentative definition of engineering-science-forensics. Engineering and forensics are about applying knowledge or discovering particulars, whereas science is the aspect of security concerned with abstracting knowledge. Justified generalization is also the key to accurate prediction.

We can tackle the problem of building general, sharable knowledge better if we see it clearly for what it is: a problem shared by all sciences, with particular strategies more or less transferable between fields depending on the details of what a field studies. Part of the solution is to integrate other fields into security, as advocated by [82]. But simply bringing in new perspectives is not enough. Morgan [65] argues that generalization, which she handles under the umbrella of resituating knowledge, is hard because knowledge is always produced locally. Knowledge transfers must be painstakingly warranted. She provides three generic strategies for this warranting process: directly local-to-local, local-to-many via abstraction followed by resituation, and constructed via exemplar representatives. Future work could fruitfully investigate how these strategies are best evidenced in security.

More specific to security, Illari and Spring [85] discuss strategies for building general, shareable knowledge used by security practitioners through examples of theft using a botnet, the intrusion kill chain, and malware reverse engineering. From these examples, they extract three challenges: the changeability of software, active adversaries, and justified secrecy among friends. These three interact to make building general, shareable knowledge particularly hard. However, such knowledge is built and shared by all three examples. Understanding explanations as mechanisms provides a good structure, though not the only one, for explaining the positive strategies used in cybersecurity. Mechanisms are clustered, both within and across fields of study, along similarities among their entities, activities, organization, etiology (history), and phenomenon of interest [85]. Further work is needed to adapt these generalization strategies to other areas of security. Another generalization strategy is mathematical modeling. Pym, Spring, and O'Hearn [75] discuss how merging logico-mathematical models and conceptual models, such as mechanistic models, improves results in program verification. Indeed, mathematical modeling, supported by both experimental and exact reasoning, can provide the same degree of support for information security engineering as it can for other, longer established forms of engineering.

It is less important to quibble over whether information security is a science than it is to lay out a satisfactory decision-making process for studying information security. It is certainly the case that information security has moved to the forefront of societal concerns. We seek to move past the debate over science or non-science. Our concern should be to identify robust decision-making and evidence-gathering tools that will allow us to make satisfactory decisions about a topic of crucial social importance. We find that none of the concerns in Section 1 are unique to security. For each, there are precedents in other scientific fields to which we can turn for advice on robust strategies.

Cybersecurity science has matured to the point that it requires specialists who reflect on the scientific enterprise itself, its own 'philosophy of' to navigate challenges of methodology, ethics, modeling choices, and interpretation of results. The importance that forensics of past events as well as building shareable, stable knowledge from particulars both play in security have few well-developed precedents in other sciences. We see great potential for future work on how to integrate these aspects of cybersecurity with modern philosophy of science. The community should move past current complaints in order to focus on questions of why the scientific process studying security may produce unsatisfactory results.

ACKNOWLEDGEMENTS

The authors would like to thank NSPW shepherds Wolter Pieters and Karl Levitt, the anonymous reviewers, all NSPW attendees, Phyllis Illari, and Marie Vasek for constructive comments on previous versions.

Spring is supported by University College London's Overseas Research Scholarship and Graduate Research Scholarship. orcid.org/0000-0001-9356-219X (Spring)

REFERENCES

[1] Anderson, R. *Security Engineering: A guide to building dependable distributed systems*, 2nd ed. Wiley, Indianapolis, IN, 2008.
[2] Anderson, R., and Moore, T. The economics of information security. *Science 314*, 5799 (2006), 610–613.
[3] Balmer, B. *Secrecy and science: A historical sociology of biological and chemical warfare.* Ashgate Publishing, Ltd., 2013.
[4] Bechtel, W., and Richardson, R. C. *Discovering complexity: Decomposition and localization as strategies in scientific research*, 1st ed. Princeton University Press, Princeton, NJ, 1993.
[5] Bickle, J. Real reduction in real neuroscience: metascience, not philosophy of science (and certainly not metaphysics!). In *Being reduced: New essays on reduction, explanation, and causation*, J. Hohwy and J. Kalestrup, Eds. Oxford University Press, 2008, pp. 34–51.
[6] Biddle, R., Chiasson, S., and Van Oorschot, P. Graphical passwords: Learning from the first twelve years. *ACM Comput. Surv. 44*, 4 (Sep 2012), 19:1–19:41.
[7] Bogen, J., and Woodward, J. Saving the phenomena. *The Philosophical Review XCVII*, 3 (1988), 303–352.
[8] Cartwright, N. *How the Laws of Physics Lie.* Clarendon Press, Oxford, 1983.
[9] Cartwright, N. Replicability, reproducibility, and robustness: Comments on Harry Collins. *History of Political Economy 23*, 1 (1991), 143–155.
[10] Cartwright, N., and Hardie, J. *Evidence-based policy: a practical guide to doing it better.* Oxford University Press, New York, 2012.
[11] Caulfield, T., and Pym, D. Improving security policy decisions with models. *IEEE Security & Privacy 13*, 5 (2015), 34–41.
[12] Collinson, M., Monahan, B., and Pym, D. *A Discipline of Math.Systems Modelling.* College Publns., 2012.
[13] Courtault, J., Galmiche, D., and Pym, D. J. A logic of separating modalities. *Theor. Comput. Sci. 637* (2016), 30–58.
[14] Craver, C. F. *Explaining the brain: mechanisms and the mosaic of unity of neuroscience.* Oxford University Press, 2007.
[15] Creath, R. Logical empiricism. In *The Stanford Encyclopedia of Philosophy*, E. N. Zalta, Ed., spring 2014 ed. Metaphysics Research Lab, Stanford University, 2014.
[16] Darden, L. *Reasoning in Biological Discoveries: Essays on Mechanisms, Interfield Relations, and Anomaly Resolution.* Cambridge University Press, 2006.

[17] Darden, L., and Maull, N. Interfield theories. *Philosophy of science 44* (1977), 43–64.

[18] Das, A., Bonneau, J., Caesar, M., Borisov, N., and Wang, X. The tangled web of password reuse. In *21st Annual Network and Distributed System Security Symposium, NDSS 2014, San Diego, California, USA, February 23-26, 2014* (2014), The Internet Society.

[19] Dawid, A. P. Beware of the DAG! *NIPS Causality: Objectives and Assessment 6* (2010), 59–86.

[20] Dear, P. *The intelligibility of nature: How science makes sense of the world.* University of Chicago Press, Chicago and London, 2006.

[21] Dittrich, D., and Kenneally, E. The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research. Tech. rep., U.S. Department of Homeland Security, Aug 2012.

[22] Dykstra, J. *Essential cybersecurity science: build, test, and evaluate secure systems.* " O'Reilly Media, Inc.", 2015.

[23] Egelman, S., Sotirakopoulos, A., Muslukhov, I., Beznosov, K., and Herley, C. Does my password go up to eleven?: The impact of password meters on password selection. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2013), CHI '13, ACM, pp. 2379–2388.

[24] Elster, J. *Nuts and bolts for the social sciences.* Cambridge Univ Press, Cambridge, UK, 1989.

[25] Evans, D., and Stolfo, S. The Science of Security: Guest editors' introduction. *Security & Privacy 9*, 3 (2011), 16–17.

[26] Evron, G. Art into Science: A conference on defense. http://artintoscience.com, Jan 25, 2017. Accessed Apr 2017.

[27] Feitelson, D. G. From repeatability to reproducibility and corroboration. *ACM SIGOPS Operating Systems Review 49*, 1 (2015), 3–11.

[28] Galison, P. Trading zone: Coordinating action and belief. *The Science Studies Reader* (1999), 137–160.

[29] Galison, P. Trading with the enemy. In *Trading zones and interactional expertise. Creating new kinds of collaboration*, M. E. Gorman, Ed. MIT Press, Cambridge, MA, 2010, ch. 3.

[30] Galison, P. Augustinian and Manichaean science. In *Symposium on the Science of Security* (National Harbor, MD, Nov 29, 2012).

[31] Galmiche, D., Méry, D., and Pym, D. The Semantics of BI and Resource Tableaux. *Math. Structures in Comput. Sci. 15* (2005), 1033–1088.

[32] Gaw, S., and Felten, E. W. Password management strategies for online accounts. In *Second Symposium on Usable Privacy and Security* (Pittsburgh, PA, USA, 2006), ACM, pp. 44–55.

[33] Geer, D. T.S. Kuhn revisited. In *NSF Secure and Trustworthy Cyberspace Principal Investigators' Meeting* (Arlington, VA, Jan 6, 2015).

[34] Given, L. M., Ed. *The Sage encyclopedia of qualitative research methods.* Sage, Thousand Oaks, CA, 2008.

[35] Glennan, S. Ephemeral mechanisms and historical explanation. *Erkenntnis 72* (2010), 251–266.

[36] Glennan, S. Mechanisms and mechanical philosophy. In *The Oxford Handbook of Philosophy of Science*, P. Humphreys, Ed. Oxford University Press, Aug 2015.

[37] Glennan, S., and Illari, P., Eds. *The Routledge Handbook of Mechanisms and Mechanical Philosophy.* Handbooks in Philosophy. Routledge, London, UK, 2017.

[38] Halpern, J. Y., and Pearl, J. Causes and explanations: A structural-model approach. Part I: Causes. *The British Journal for the Philosophy of Science 56*, 4 (2005), 843–887.

[39] Halpern, J. Y., and Pearl, J. Causes and explanations: A structural-model approach. Part II: Explanations. *The British Journal for the Philosophy of Science 56*, 4 (2005), 889–911.

[40] Hatleback, E., and Spring, J. M. Exploring a mechanistic approach to experimentation in computing. *Philosophy & Technology 27*, 3 (2014), 441–459.

[41] Hatleback, E. N. The protoscience of cybersecurity. *The Journal of Defense Modeling and Simulation* (2017), 1–8.

[42] Hempel, C. G. The function of general laws in history. *Journal of Philosophy 39* (1942), 35–48.

[43] Herley, C., and van Oorschot, P. SoK: Science, security, and the elusive goal of security as a scientific pursuit. In *Symposium on Security and Privacy (Oakland)* (San Jose, CA, May 22-24, 2017), IEEE.

[44] Illari, P. M., and Williamson, J. What is a mechanism? Thinking about mechanisms across the sciences. *European Journal for Philosophy of Science 2*, 1 (2012), 119–135.

[45] Ishtiaq, S. S., and O'Hearn, P. W. BI as an assertion language for mutable data structures. In *Principles of Programming Languages* (London, UK, 2001), ACM, pp. 14–26.

[46] Jain, R. *The Art of Computer Systems Performance Analysis.* Wiley & Sons, 1991.

[47] Jasanoff, S. *The fifth branch: Science advisers as policymakers.* Harvard University Press, Cambridge, MA, USA, 1990.

[48] Katz, J. Call for papers: Hot topics in the science of security (HoTSoS), Dec 2016. http://cps-vo.org/group/hotsos/cfp.

[49] Killourhy, K., and Maxion, R. Comparing anomaly-detection algorithms for keystroke dynamics. In *Dependable Systems & Networks* (Lisbon, Portugal, Jun 2009), IEEE, pp. 125–134.

[50] Koen, B. V. *Discussion of the method: Conducting the engineer's approach to problem solving.* Oxford University Press, New York, 2003.

[51] Kott, A. Towards fundamental science of cyber security. In *Network Science and Cybersecurity*, R. E. Pino, Ed. Springer, New York, NY, 2014, pp. 1–13.

[52] Krol, K., Spring, J. M., Parkin, S., and Sasse, M. A. Towards robust experimental design for user studies in security and privacy. In *Learning from Authoritative Security Experiment Results (LASER)* (San Jose, CA, 2016), IEEE, pp. 21–31.

[53] Kuhlmann, D., Chen, L., and Mitchell, C. Trust and legitimacy in security standardization – a new management issue? In *Interoperability for Enterprise Systems and Applications (I-ESA 16)* (Guimaraes, Portugal, Mar 29-Apr 1 2016), ISTE Publications.

[54] Kuhn, T. S. *The structure of scientific revolutions*, 4th ed. University of Chicago Press, Chicago and London, 2012. Introductory essay by Ian Hacking.

[55] Leonelli, S. *Understanding in biology: The impure nature of biological knowledge.* University of Pittsburgh Press, Pittsburgh, PA, USA, 2009, pp. 189–209.

[56] Machamer, P., Darden, L., and Craver, C. F. Thinking about mechanisms. *Philosophy of science 67* (March 2000), 1–25.

[57] Maxion, R. Structure as an aid to good science. In *Workshop on the Science of Cyber Security* (Bristol, UK, January 2015), IFIP Working Group 10.4.

[58] Meijers, A., Ed. *Philosophy of Technology and Engineering Sciences*, vol. 9 of *Handbook of the Philosophy of Science.* North-Holland, Amsterdam, 2009.

[59] Metcalf, L., and Casey, W. *Cybersecurity and Applied Mathematics.* Syngress, Cambridge, MA, USA, 2016.

[60] Meushaw, R., Ed. *Developing a blueprint for a science of cybersecurity* (Fort Meade, MD, 2012), vol. 19:2 of *The Next Wave*, U.S. National Security Agency.

[61] Meushaw, R. What is security science?, Oct 19, 2012. http://cps-vo.org/node/6041.

[62] Mitchell, S. D. *Biological complexity and integrative pluralism.* Cambridge University Press, Cambridge, 2003.

[63] MITRE Corporation. Science of cyber-security. Tech. Rep. JSR-10-102, JASON Office, McLean, VA, Nov 19, 2010.

[64] Morgan, M. S. Nature's experiments and natural experiments in the social sciences. *Philosophy of the Social Sciences 43*, 3 (2013), 341–357.

[65] Morgan, M. S. Resituating knowledge: Generic strategies and case studies. *Philosophy of Science 81*, 5 (2014), 1012–1024.

[66] Nagel, E. *The structure of science: Problems in the logic of scientific explanation*, 2nd ed. Routledge & Kegan Paul, London, 1979.

[67] National Cyber Security Centre. Password guidance: Simplifying your approach, 2017. https://www.ncsc.gov.uk/guidance/password-guidance-simplifying-your-approach.

[68] National Science Foundation. Federal Cyber Service: Scholarship for Service (SFS). A federal cyber service training and education initiative. Tech. Rep. NSF 01-167, NSF, Directorate for education and human resources, Division of undergraduate education, Arlington, VA, 2001.

[69] Norton, J. D. There are no universal rules for induction. *Philosophy of Science 77*, 5 (December 2010), pp. 765–777.

[70] O'Hearn, P., and Pym, D. The logic of bunched implications. *Bulletin of Symbolic Logic 5(2)* (1999), 215–244.

[71] Oltramari, A., Cranor, L. F., Walls, R. J., and McDaniel, P. D. Building an ontology of cyber security. In *Semantic Technology for Intelligence, Defense, and Security* (Fairfax, VA, USA, Nov 2014), pp. 54–61.

[72] Oram, A., and Wilson, G. *Making software: What really works, and why we believe it.* O'Reilly Media, Inc., 2010.

[73] Pearl, J. *Causality.* Cambridge University Press, Cambridge, UK, 2009.

[74] Popper, K. R. *The logic of scientific discovery.* Hutchinson, London, 1959.

[75] Pym, D., Spring, J., and O'Hearn, P. Why separation logic works. Submitted, 2017. Manuscript: http://www0.cs.ucl.ac.uk/staff/D.Pym/PSO-SL.pdf.

[76] Research Institute in Science of Cyber Security. Annual report. Tech. rep., University College London, London, UK, 2016.

[77] Reynolds, J. C. Separation logic: A logic for shared mutable data structures. In *Proceedings of the 17th Annual IEEE Symposium on Logic in Computer Science* (Washington, DC, USA, 2002), LICS '02, IEEE Computer Society, pp. 55–74.

[78] Royal Society. Progress and research in cybersecurity: Supporting a resilient and trustworthy system for the UK. Tech. Rep. ISBN: 978-1-78252-215-7, London, UK, July 2016.

[79] Scarfone, K., and Mell, P. Guide to intrusion detection and prevention systems. Tech. Rep. SP 800-94, U.S. National Institute of Standards and Technology, Gaithersburg, MD, Feb 2007.

[80] Shannon, G., Bogner, K., Epstein, J., Fraser, T., King, S., Martin, W. B., Maughan, D., Morrow, J., Newhouse, W., Polk, W. T., and Vagoun, T. Federal cybersecurity research and development strategic plan: Ensuring prosperity and national security. Tech. rep., National Science and Technology Council, Washington, DC, Feb 2016.

[81] Shirey, R. Internet Security Glossary, Version 2. RFC 4949 (Informational), Aug. 2007.

[82] Shostack, A., and Stewart, A. *The new school of information security.* Pearson Education, 2008.

[83] Simon, H. A. *The sciences of the artificial*, 3rd ed. MIT press, Cambridge, MA,

1996.

[84] Spring, J. M., and Hatleback, E. Thinking about intrusion kill chains as mechanisms. *Journal of Cybersecurity 2*, 2 (2017).

[85] Spring, J. M., and Illari, P. Mechanisms and generality in information security. *Under review* (2017).

[86] Stake, R. E. *The art of case study research.* Sage, Thousand Oaks, CA, 1995.

[87] Stodden, V. Reproducing statistical results. *Annual Review of Statistics and Its Application 2* (2015), 1–19.

[88] The Economist. The city of the century: How Vienna produced ideas that shaped the west, Dec 24, 2016.

[89] Uebel, T. Vienna circle. In *The Stanford Encyclopedia of Philosophy*, E. N. Zalta, Ed., spring 2016 ed. Metaphysics Research Lab, Stanford University, 2016.

[90] University College London. The Research Institute in Science of Cyber Security (RISCS). https://www.riscs.org.uk/, 2017. Accessed Mar 6, 2017.

[91] Ur, B., Kelley, P. G., Komanduri, S., Lee, J., Maass, M., Mazurek, M. L., Passaro, T., Shay, R., Vidas, T., Bauer, L., Christin, N., and Cranor, L. F. How does your password measure up? The effect of strength meters on password creation. In *USENIX Conference on Security Symposium* (Bellevue, WA, 2012), USENIX Association, pp. 65–80.

[92] Vincenti, W. G. *What engineers know and how they know it: Analytical studies from aeronautical history.* Johns Hopkins Studies in the History of Technlogy. Johns Hopkins University Press, Baltimore and London, 1990.

[93] Wash, R. Folk models of home computer security. In *Proceedings of the Sixth Symposium on Usable Privacy and Security* (New York, NY, USA, 2010), SOUPS '10, ACM, pp. 11:1–11:16.

[94] Williamson, J. Evaluating evidence in medicine. https://blogs.kent.ac.uk/jonw/projects/evaluating-evidence-in-medicine/, Jun 1, 2015.

[95] Winn, J. K. Should vulnerability be actionable? Improving critical infrastructure Computer security with trade practices law. *George Mason Univ. Critical Infrastructure Protection Project Papers Vol. II* (2004).

[96] Woodward, J. *Making things happen: A theory of causal explanation.* Oxford University Press, Oxford, UK, 2003.