

- . The Science of Security, and the Future
- . "Does a field make progress because it is a science,
- . or is it a science because it makes progress?"
- . Dan Geer, RSA, 23 April 15

Predicting the future is, as you know, fraught with difficulty. I have no better handle on the detailed future than you do because we both have nearly zero real predictive power. We can, however, look at what the drivers of the future are. One is culture, the other is science. I will focus on science today, but on culture let me quote Jack Bogle, the founder of Vanguard:

"In recent years, annual trading in stocks ... averaged some \$33 trillion. But capital formation -- that is, directing fresh investment capital to its highest and best uses, such as new businesses, new technology, medical breakthroughs, and modern plant and equipment for existing business -- averaged some \$250 billion. Put another way, speculation represented about 99.2 percent of the activities of our equity market system, with capital formation accounting for 0.8 percent." [JB]

Short-sighted-ness is what Bogle is talking about, and, dare I say, the cybersecurity field is second only to finance in that and that precisely. I'll take up the art of the long view another day.

For today and for the science of security, I will re-visit T.S. Kuhn's landmark work, The Structure of Scientific Revolutions. I rather suspect the few of you who have read it did so as an assigned reading in some classroom long ago. It was published in 1962 by the University of Chicago as a volume in the International Encyclopaedia of Unified Science, a project that was never, in fact, completed.

I am not a scientist and do not think of myself as one. As such, I will attempt to talk about science but I disclaim that I've earned any right to do so.

Today, the science of security is advanced in a number of ways, but in the U.S. the primary investor is the National Science Foundation. I am, myself, 37 days older than the NSF, and I have a deep respect for what PL 81-507 created the NSF to do, which in full is this:

To promote the progress of science; to advance the national health, prosperity, and welfare; to secure the national defense; and other purposes.

Kuhn's book, which he consistently refers to as an essay, is basically about what science is, based on the observables of what science does and has done. Between him, his supporters and his critics, many noted philosophers, and others, what science *is* is to this day unsettled at its core. Perhaps that is why both Kuhn's supporters and his critics agree on one thing: There is no algorithm to science, and without an algorithm, prediction is more about luck than deduction.

I come at this question, the question of what is science and, later in this essay, whether cybersecurity can or has become a science, with a set of biases formed by how I was myself taught. For the record, I am the son of a Certified Public Accountant and was myself formally trained first as an electrical engineer and then further

as a biostatistician. To the extent that one can assess one's own biases, mine are that while numbers are endlessly fascinating, they are not to be sought for their own sake but rather for the purpose of decision support. Medicine, where I worked for fifteen years, has a term of art that encapsulates that idea; the term is "no therapeutic difference" meaning that if a blood test or a scan or what have you might improve the subtle precision of a subtle diagnosis, if the therapy would nevertheless remain the same, then the pursuit of further diagnostic precision is not medicine but something else. As we say on the security metrics mailing list, I am a measurer, not a modeler.

Biases aside, the time has come to read Kuhn's essay in the context of cybersecurity. He begins and ends with what is a circular idea, that a scientific community is defined by what beliefs practitioners share, and what beliefs practitioners share defines what community they are in. This is, in fact, instructive as no science begins in mature form, but rather any new science will begin in much more modest circumstances where, in fact, there is nothing approaching a consensus in any sense of the word, that, early on, consensus is not even a concept. As such, part of becoming a mature science is the development of a broad consensus about the core concerns of that branch of knowledge.

Kuhn's word for the collections of exemplars of good science was "paradigm," a word whose meaning today is all but entirely Kuhn's, even amongst those who never read a blessed word he wrote. Anyone from Berkeley or Madison or Cambridge will remember the bumper sticker "subvert the dominant paradigm" and thus attest to the impact of Kuhn's propositions bundled up in that one word, even on anarchists.

But what is a "paradigm" and why do we want one? As Kuhn puts it, "[Paradigms] are the source of the methods, the problem field, and the standards of solution accepted by any mature scientific community at any given time." Kuhn's book and the two decade long back and forth between Kuhn and philosophers notwithstanding, the simplest version is that a paradigm is all the things that a scientist can assume that his or her colleagues will congenially understand about their common work without explicitly explaining them or arguing them from first principles again and again.

Several authors besides Kuhn have adopted that idea to explain the proliferation of jargon as being something that markedly improves the efficiency of communication between practitioners who do, in fact, share a paradigm, who share a set of techniques, trainings, and world view in common. As Kuhn said, "Although it is customary, and is surely proper to deplore the widening gulf that separates the professional scientist from his colleagues in other fields, too little attention is paid to the essential relationship between that gulf and the mechanisms intrinsic to scientific advance." In other words, impenetrable jargon between scientists sharing a paradigm is a side-effect of that sharing, and a tool of efficiency because it insulates science from society.

But, I hear you say, any field from poetry to American history to macroeconomics can have an impenetrable jargon, does that make those fields sciences? Clearly not, and here is where Kuhn's paradigm construct is central. In a science, the shared paradigm is universal whereas in the humanities, say, there are always competing schools of thought that are unlikely ever to sign some unifying intellectual

peace treaty. As Kuhn would have it, a paradigm -- a shared framework for how the world works -- is the engine for creating the kinds of puzzles that individual scientists are able to solve if they work hard enough, the paradigm creates general agreement amongst those who hold it as to where further research is needed. When a body of scientists is asked "Where should research go next?" the answer will be the same as to the question "Where is research going?" because the shared paradigm makes it so. And note that I used Kuhn's word, "puzzles," rather than "problems." Kuhn is at some pains to make clear that a scientist doing science is solving the puzzles that the paradigm leaves open to solution. He or she is not solving problems in the vernacular sense of the word "problem."

Again quoting Kuhn, "Men whose research is based on shared paradigms are committed to the same rules and standards for scientific practice. That commitment and the apparent consensus it produces are prerequisites for normal science, i.e., for the genesis and continuation of a particular research tradition... Acquisition of a paradigm and of the more esoteric type of research it permits is a sign of maturity in the development of any given scientific field."

Science is about knowledge derived from experiment whether those experiments are designed or natural, and the refinement of knowledge is to be ultimately found in quantification, but though the ability to explicitly quantify is a sign of agreement on terms and processes, the ability to quantify is not itself proof that the body of knowledge involved is of a scientific sort -- the ability to quantify is necessary but not sufficient. To be blunt, no less a personage than Nassim Taleb[NNT] said this of economics: "You can disguise [its] charlatanism under the weight of equations, and nobody can catch you since there is no such thing as a controlled experiment."

Competing schools of thought are, however, always present before a mature science first appears. As Kuhn said, "What is surprising, and perhaps also unique in its degree to the fields we call science, is that such initial competing schools should ever largely disappear. For they do disappear to a very considerable extent and then apparently once and for all." Perhaps it thus possible to say that topics of study that never coalesce their competing schools are either fated never to be sciences or are in some state of arrested development that may someday be cured. Those that can and will, but have not yet done so, are what Kuhn called pre-paradigmatic, meaning not yet a science. The appearance of a paradigm that all can accept is when the transition to being a science occurs, or, as Kuhn put it, "Except with the advantage of hindsight, it is hard to find another criterion that so clearly proclaims a field a science."

For Kuhn, the appearance of a paradigm transforms those who merely study first into a discipline and then into a profession, even if as the science develops its thought processes its language become ever less intelligible to non-specialists. One can even say, and Kuhn does, that the paradigm itself is the last result of the science in question that can be appreciated by the lay audience -- after that all progress is in journal articles not readable by non-specialists, enough so that "The scientist who writes [for the lay reader] is more likely to find his professional reputation impaired than advanced" for having done so.

Now as everyone here knows, from time to time a science may undergo a revolution, which in Kuhn's terms is precisely the laying down

of one paradigm in preference for another. The title of his book is to be understood as precisely that, that scientific revolutions share aspects of structure that we can now describe as there have been enough of them in the last 400 years to discern that structure. If you consider physics to be the paragon of a hard science, then the transition from Newtonian mechanics to Einsteinian relativity demonstrates exactly the point Kuhn was making, that there comes a moment when research has reached a kind of impasse where the nature of what now look to be puzzles needing further study cannot be profitably investigated within the paradigm that now holds.

Kuhn referred to these impasses as the appearance of an anomaly, one that the existing paradigm cannot evaluate by way of further research consistent with the paradigm then in place. His review of past revolutions centered in each case on the appearance of irreconcilable anomalies that made a given field ripe for revolution. That roasting metals caused them to gain weight thus indicating that they had absorbed some fraction of the air around them, a fraction that could be exhausted, led to the idea that air might not be the one and only gas but rather a combination of gases. Perhaps more significantly to the very idea of revolution is that even though Lavoisier had discovered oxygen, others in the field, notably Priestly, never accepted the existence of oxygen and held to the phlogiston theory to the end of their careers. I say "more significantly" as the trite version of "What is a scientific revolution?" is that it is a time when newcomers to the field adopt the new paradigm while those already in the field slowly die off. It is a generational change.

The first and perhaps main objection to Kuhn's thesis was that it implied a certain irrationality to the advancement of science, that science and scientists did not, despite what textbooks tend to say, advance in a monotonic way wholly under the self-imposed discipline of rationality. Kuhn's implication that scientific advance is not necessarily rational but may occasionally be irrational was found by many to be unacceptable. I'm not going to argue that point here as this essay is not about the history of science, but I will quote George Bernard Shaw, "The reasonable man adapts himself to the world; the unreasonable one persists in trying to adapt the world to himself. Therefore all progress depends on the unreasonable man."

This is not to say that a newer paradigm must be uniformly better than its predecessor, only that it account for the heretofore irreconcilable anomalies that scientists have found within the older paradigm. Ptolemaic earth-centric astronomy explained the motion of the planets fairly well, but its complexity was increasing far more rapidly than its accuracy, thus a crisis and a ripeness for a new paradigm. Kuhn observed that a "proliferation of competing theories [is] the concomitant of crisis" and it is the new paradigm that resolves the crisis. He goes on to say that "So long as the tools a paradigm supplies continue to prove capable of solving the problems it defines, science moves fastest and penetrates most deeply through confident employment of those tools. The reason is clear. As in manufacture so in science -- retooling is an extravagance to be reserved for the occasion that demands it. The significance of crises is the indication they provide that an occasion for retooling has arrived."

Kuhn's idea of crisis is the dual of his idea of paradigm. Where a paradigm suggests puzzles that further research will solve, in a

crisis this is not so. Yet the occasional crisis is itself necessary for advancement as any paradigm whose theories completely explain all observable fact ceases to be science and becomes engineering. In other words, the crisis is not the end of research but the substitution of a new paradigm for an old and a new set of research puzzles awaiting solution. Just as "a scientific theory is declared invalid only if an alternate candidate is available to take its place," "to reject one paradigm without simultaneously substituting another is to reject science itself." In short, when an anomaly appears, there are only three resolutions available: (1) solve the problem, (2) leave the problem for future scientists, or (3) use the crisis to force a new paradigm on the field. One can almost hear Rahm Emmanuel's political dictum to "Never let a good crisis go to waste" in parallel to Kuhn.

When there is a shift of paradigm, that is to say a scientific revolution, it may serve to redirect a field so completely that some parts of it fall away entirely, the separation of astronomy from astrology or the separation of chemistry from physics to pick two examples. As Kuhn put it, the choice between paradigms is a choice between incompatible modes of community life. And, yes, he does mean "scientific community" in an altogether social sense; "Since no two paradigms leave all the same problems unsolved, paradigm debates involve the question: which problems is it more significant to have solved?"

Kuhn also views the education and apprenticeship of a new scientist as the propagation of the field's paradigms to the new scientist such that he or she can join the professional scientific community. Here he argues with how that is done in saying that "Textbooks refer only to that part of the work of past scientists that can easily be viewed as contributions to the statements and solutions of the text's paradigmatic problems." In other words, the supposed linearity and unbroken advance in a field that is the typical narrative of textbook writing is wrong, rather it is that science is cumulative within the sway of a particular paradigm, but only then. "Scientists are not, of course, the only group that tends to see its discipline's past developing linearly toward its present vantage. The temptation to write history backward is both omnipresent and perennial. But scientists are more affected by the temptation to rewrite history, partly because the results of scientific research show no obvious dependence upon the historical context of the inquiry, and partly because, except during crisis and revolution, the scientist's contemporary position seems so secure." Alfred North Whitehead said as much, "A science that hesitates to forget its founders is lost" though Kuhn goes to some length to call that process "reinterpretation" rather than "forgetting," though with the same result, namely that scientific progress is not linear even if every textbook speaks of standing on the shoulders of giants.

It is tempting to analogize science as a kind of biologic process. The noted paleobiologist, Stephen Jay Gould, famously coined the term of "punctuated equilibrium" to describe the fossil record, saying that there would naturally be long periods of stasis followed by much shorter periods of rapid change. He wrote "The history of life is not a continuum of development, but a record punctuated by brief, sometimes geologically instantaneous, episodes of mass extinction and subsequent diversification." [SJK] Extending Gould, Matthew Stoneking [MS] considers scientific revolutions to be quite similar to extinction events and, more importantly, the diversifying speciation that follows an extinction event. In this, Stoneking

and Kuhn diverge in that Stoneking sees scientific speciation as a consequence of revolution where Kuhn sees proliferating species, that is to say theories, as the precipitating event just prior to a revolution.

Stoneking goes on to say "If all members of a species were identical in all features, then the species would be extremely vulnerable to even minor changes in the environment. As it is, when the environment changes, there are usually members of the species who are better suited to survive under the new conditions. The offspring of those individuals tend to be more numerous and better able to compete for resources than the offspring of other members of the species. The analogy to periods of normal scientific activity are crude, but fit pretty well. During periods of normal science, members of a given discipline adapt the paradigm theory to fit observation in as many specific cases as possible. The outcome of that work is a better articulated paradigm, one that works better in a wider range phenomena. During periods of normal science, the paradigm evolves in a smooth way, in a way analogous to the minor adaptations of species to short term changes in the environment."

In other words, Stoneking is arguing that Gould and Kuhn are of the same view, that scientific communities and scientific disciplines are species, that linguistic isolation within such communities serve the purpose of reproductive isolation, that interdisciplinary fields have the characteristics of biologic hybrids both as to vigor and to reproductive infidelity, that a scientific community in crisis corresponds to an impending extinction event, and that scientific revolutions are speciation events. I think he pushes that a bit too far, as, with Kuhn, I don't see a scientific revolution as a speciation event but rather speciation events as the prodrome of an extinction event, that is to say an event where survival of the fittest is coldly demonstrated. What Kuhn does see is this, "Though science surely grows in depth, it may not grow in breadth as well. If it does, that breadth is manifest mainly in the proliferation of scientific specialties, not in the scope of any single specialty alone."

And to round out Kuhn, the single most controversial part of his analysis is that science is not, despite pretensions to the contrary, unperturbed by the social milieu in which it operates. For Kuhn, science cannot be the search for ultimate truth, which he makes clear when he says "We may have to relinquish the notion, explicit or implicit, that changes of paradigm carry scientists and those who learn from them closer and closer to the truth." It is his conclusion that science is undirected that offends the sensibilities of many scientists. Kuhn's discussion of aesthetic qualities such as theoretic elegance elaborates this further, which Brian Greene reinforces, writing that "It is certainly the case that some decisions made by theoretical physicists are founded upon an aesthetic sense -- a sense of which theories have an elegance and beauty of structure on a par with the world we experience. Of course, nothing ensures that this strategy leads to truth." [BG]

I would like now to turn toward cybersecurity. I use the term "cybersecurity" in lieu of the many available alternatives, perhaps because I work alongside the governmental policy community where the word cybersecurity resonates and has become a term of both art and of funding. If you prefer "information security" or "computer security" or something else again, please hear those phrases and don't let my choice of terms derail the rest of this essay.

One of the first questions we might ask is whether cybersecurity is a science or, if not, whether it ever will be. With several others, I am one of the expert reviewers for the National Security Agency's annual "Science of Security" competition and award.[SOS] Quoting its rationale, "The competition was established to recognize the current security paper that best reflects the conduct of good science in the work described. [Science of Security] is a broad enterprise, involving both theoretical and empirical work. While there can only be one best paper, any one paper cannot span that full breadth. Nonetheless, the field is broad and work in all facets is encouraged and needed. The common denominator across the variety of approaches is solid methodology and effective communication, so those aspects of the papers [are] strong factors in our decision."

Papers are nominated for consideration, and I encourage you to do so, but I am also here to report that amongst the reviewers our views of what constitutes a, or the, Science of Security vary rather a lot. Some of us would prioritize purpose, agreeing with Charles Darwin that "All observation must be for or against some view if it is to be of any service." [CD] 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. We vary, and I take that to be a vote of sorts on whether cybersecurity is yet a science.

The question of whether cybersecurity is yet a science is a hard one. I am sorely tempted to answer the question "Is cybersecurity a science" with "Getting closer, but not yet" -- to say, in other words, that we are in the pre-paradigmatic stage with a variety of schools of thought. I'll return to that later, but let's talk for a bit about candidate paradigms of cybersecurity. If they exist and have turned over from time to time, then my answer would be simply wrong. But let me be clear about one thing that may make cybersecurity different than all else and that is that we have sentient opponents. The physicist does not. The chemist does not. Not even the economist has sentient opponents. We do. What puzzles we have to solve are not drawn from some generally diminishing store of unsolved puzzles, nor could our theories completely explain all observable fact thus reducing our worries and our work to engineering alone. There is something different about a search for truth when there isn't any, or at least any that lasts long enough to exhaustively explore.

And while I view the greatest potential contradiction between cybersecurity and science to be the sentient opponent, it may be that it is the rate of change of the technical fabric which is ultimately distinguishing. Perhaps rate of change is a continuum and when I say that the sentient opponent is the greatest issue it may just be that the gross rate of change is the sum of technical advance and sentient opponents with sentient opponents being the most accelerating fraction of the overall sum we call technologic change.

Take one of the most basic tools we employ, that of authentication. Authentication is the solution to the puzzle of identity establishment, a puzzle that derived from the paradigm of perimeter control. For authorization to have meaning, authentication had to come first, and authentication is, of necessity, established prior to its use

in any authorization decision. But this mechanism is the solution to a puzzle consequent to the paradigm in which it has meaning, the paradigm of perimeter control. We came to the idea of perimeter from the physical world where walled cities date almost to the time that there were cities at all. Even defense in depth is hardly new; the concentrically walled Irish fort known as Dun Aengus is at least 3000 years old. With the technologic change that may as well be called the Internet, it is only natural that we transposed the paradigm of perimeter control to the digital world.

The paradigm of perimeter control has been in an evident crisis for some time now. The crisis is not merely because the definition of perimeter may have been poorly applied in practice, but because some combination of always-on and universal addressability collectively make the paradigm of a defensible perimeter less and less a paradigm where research is itself likely to patch up the mess and retain the core and guiding paradigm of perimeter control. The parallel with Ptolemaic astronomy is pretty fair. On the one hand, every improvement in observational accuracy made the motions of the planets more complicated to describe with epicycles upon epicycles. On the other hand, our hand, the threat to systems from always-on universal addressability has become too rich to be just a new set of puzzles solely within the paradigm of perimeter control -- the defensible perimeter began to have its own version of epicycles within epicycles by a shrinking of what a perimeter could, or should, control.[DG]

A second crisis for the paradigm of perimeter control is upon us now and that is perhaps best exemplified with a commercial example. Let's count cores in the Qualcomm Snapdragon 801. The central CPU is 4 Cores, the Adreno 330 GPU another 4, Video Out is 1 more, the Hexagon QDSP is 3, the Modem is at least 2 and most likely 4, Bluetooth is another 1 as is the USB controller and the GPS. The Wifi is at least 1 and most likely 2, and none of this includes charging, power, or display. That makes somewhere between 18 and 21 cores. In the vocabulary of the Internet of Things, I ask you whether that is one thing or the better part of two dozen things? It is pretty certain that each of those cores can reach the others, so is the perimeter to be defended the physical artefact in the user's pocket or is it the execution space of each of those cores?

I looked at seven different estimates of the growth of the Internet of Things as a market phenomenon, everything from smart electric meters to networked light bulbs to luxury automobiles and the median is a compound annual growth rate of 35%. If perimeter control is to remain the paradigm of cybersecurity, then the number of perimeters to defend in the Internet of Things is doubling every 17 months.

If the paradigm of perimeter control is no longer producing puzzles that can be solved by further scientific research, then what? Noting, as Kuhn does, that "to reject one paradigm without simultaneously substituting another is to reject science itself," what might be a substitution? While that is for you, the security community, to decide, I'll say that one of the alternates is a mix of surveillance and accountability. If, as seems to be the case, everything we are or do is unique if examined closely enough, then the idea of authentication as verifying an assertion like "My name is Dan" can easily morph into an observable like "Sensors say that this is Dan." In other words, our paradigm of an authentication transaction before any other perimeter piercing transaction is itself showing its age. Part of the authentication crisis is a fork in the road that cannot be avoided here, and that is that the

more authorities you personally command, the more varieties of authentication you have to have or, at least, you have had to have up until now.

The paradigm that is the obvious alternative to perimeter control and thus authentication as a gating function is accountability based on one single unspoofable identity per person. If I am right, that real soon now identity is simply an observable that needs no assertions, then that single identity which the individual has but does not need to prove may be fast upon us. The National Strategy for Trusted Identities in Cyberspace is not worded in that way though that is how I read it, but, in fairness, the deliberative writing of national strategies in the face of accelerating change is as hard as providing cybersecurity in the face of accelerating change.

Nevertheless, if being part of the modern world in no more robust way than appearing unmasked on a public street is the same as submitting to unitary identity observable at a distance by things you never heard of, then for the individual that means either submission or withdrawal. The individual's choice is outside a talk on the science of security, unless in the term "security" you want to include the latent power of data collected for no particular purpose by an Internet of Things growing 35% per annum.

Note that Kuhn never said that a switch to a new paradigm would be delightful or comforting, he merely said that it would better explain the way the world works while suggesting new puzzles for scientists who share that paradigm to pursue. Authentication transactions as a prodrome to authorization transactions in service to perimeter control may soon be behind us, including in the peer to peer world. If, in fact, being authenticated as yourself is unavoidable, then there is no proving that this is the Dan for whom there is a book entry allowing him into some robot-protected building but rather an accountability regime based on whether that Dan did or did not enter a building for which he might later be penalized. His crime would not be masquerading as some identity other than his own so as to get in, but rather that of he was observed to have gone in even though he was forbidden to do so.

Let's try another all but equivalent paradigm, namely that of personal control of personal data. You can argue that it is little different than the cybersecurity paradigm of perimeter control in the large, but now confined only to the small. I disagree; if the person is to have free will, then the person has to have, as Eric Hughes put it so long ago, "...the power to selectively reveal oneself to the world." [EH] This is why in my own work I have defined a state of privacy as whether or not you retain the effective capacity to misrepresent yourself. [DG2]

But the paradigm that Eric and I shared is now in substantial crisis because of metadata collection both by government agencies and by advertising agencies. There is no mechanistic difference whatsoever between personalization and targeting save for intent of the analyst. Daniel Solove's work notwithstanding, [DS] the meme of "I have nothing to hide" has captured the field. With self-driving cars and electronic health records and smart(er) electric grids, there are unprecedented advantages to the individual and to society. There is also a crisis of paradigm.

In this case, the paradigm is that data to be shared has to be

specifically permitted in discrete chunks. Nothing trendy coming to market from the kinds of startups I see shares that paradigm. Nothing coming from regulatory agencies shares that paradigm except in vague idiosyncracies. The paradigm of specific permission and discrete chunks is in crisis. It is as if the paradigm has been rejected and a new one adopted based on hope alone.

We have known for some time that traffic analysis is more powerful than content analysis. If I know everything about to whom you communicate including when, where, with what inter-message latency and at what length, from what geolocation and with what device, then I know you. If all I have is the undated, unaddressed text of your messages, then I am an archaeologist, not a case officer. The soothing mendacity of proxies for the President who said "It's only metadata" was to rely on the ignorance of the listener. But you, here, know all that.

What I am suggesting as the crisis around the paradigm of selective revelation is that as with metadata, there is so much redundancy in what is observable that prohibiting one or another form of collection has no meaningful effect whatsoever on those agencies, intelligence or advertising, who would build a model of you from metadata alone. As but one example, with current technology I can read the unique radio signature of your beating heart at five meters. As with anything that has an electromagnetic output, the only technologic question is the quality of the antenna. If I can take your picture on the public street without your permission or notice, why can't I record your heart? Or your iris? Or your gait? Or the difference in temperature between your face and your hands? That list is long and getting longer. It is a crisis for which the paradigm of selective revelation can scarce put up puzzles fast enough, and scientific solving of those puzzles can, at best, trail the curve.

So what might be an alternate paradigm, one that can replace the paradigm of selective revelation as a shared world view and source of research puzzles worth solving for privacy scientists? I will suggest one, and it goes like this: Putting aside, for the moment, questions of morality, if the citizenry of a democracy choose a path, then that path cannot be wrong, it can only be real. Our citizenry has chosen a technopolitical framework that involves everything from wearable health monitors to self-driving cars to Internet-connected thermostats to Lojack for children and so forth. Again, that cannot be wrong, it merely is. Because the public chooses in that way, the world now admits a set of problems for which science had better find a generative paradigm.

Perhaps what heretofore we have known as confidentiality is becoming quaint. And irrelevant. Perhaps science will have to reposition confidentiality within some new paradigm that prioritizes integrity, not confidentiality. Perhaps a world in which data can and will be collected irrespective of selective permission granting is a world in which the data had better be right. If more and more intelligent actors are to be out there doing our implicit bidding long after we've forgotten their configuration interface, then data integrity had better be as absolute as we can make it, and that is then where the research puzzles will have to be found.

If we are to have all-electronic health records and regular monitoring by everything from our toilet to the breathalyzer in our cars all the while the the majority of medicines transition to being genomically

personalized, we had better be sure that data integrity is paradigmatic. The longstanding triad of confidentiality, integrity, and availability may now be contracted to integrity and availability. Only this past October, the Santa Fe Institute and Morgan Stanley held a joint symposium entitled "Are Optimization and Efficiency the Enemies of Robustness and Resilience?"[SFI] so perhaps the crisis is already in the early stage of being made clear.

Perhaps I have it wrong, perhaps the topmost paradigm of the science of security is simply that of defense. Perhaps the rise of sentient opponents makes that paradigm of defense unarguable. Perhaps that is the paradigm, as evidenced by rafts of paradigmatically generated puzzles of the sort of how can this or that be hardened or otherwise defended, up to and including DARPA's Grand Challenge where a "capture the flag" contest will be entirely robotic.[DARPA]

If defense is and has been our paradigm, then that, too, is in crisis. That is in no way a failure; paradigms only change due to the success that the one paradigm has in motivating science to explore the world thoroughly enough to discover anomalies that cannot be made to fit within the paradigm that caused them to be discovered in the first place. The outgrowth of the paradigm of defense has been guidance that has allowed us, including non-scientist practitioners, to get better and better. We have discovered and then deployed better tools, we have come to understand causal chains and thus to better understood practices, and by way of the educational byproduct of such researches we have more, and better, colleagues. That's the plus side, and it is one terrific plus side. But if I am interested in the ratio of skill to challenge, then, as far as I can estimate, we are expanding the society-wide attack surface faster than our science of security is expanding our collection of tools, practices, and colleagues. If your island nation is growing more and better food, that's great. If your population is growing faster than those improvements in food production can keep up, that's bad. Society's adoption curves for new technology look ever-steeper from where I sit. The paradigm of defense is in crisis.

Part of my feeling stems from a long-held and well-substantiated belief that all cybersecurity technology is dual use. Perhaps dual use is a truism for any and all knowledge, as well as any and all tools derived from knowledge, that knowledge can be used for good or ill -- but I am convinced that dual use is inherent in cybersecurity tools. If your definition of "tool" is wide enough, I suggest that the cyber security tool-set strongly favors offense these days. Chris Inglis, recently retired NSA Deputy Director, remarked that if we were to score cyber the way we score soccer, the tally would be 462-456 twenty minutes into the game, i.e., all offense. I will take his comment as confirming at the highest level not only the dual use nature of cybersecurity but also confirming that offense is where the innovations that Nation States, and only Nation States, can afford is going on. Does that not change the cybersecurity paradigm? Even if you are not willing to trade in the paradigm of defense for a paradigm of offense, you will at least need to modify the paradigm of defense to something like a paradigm of dual-use-cognizant defense.

One embodiment of the paradigm of defense has been the movement to build security in. The successes of that movement are precisely of the sort I mentioned before when I said that we have discovered and then deployed better tools, we have come to understand causal chains and thus to better understood practices, and by way of the

educational byproduct of such researches we have more, and better, colleagues. But to remind you of the truism in Adi Shamir's 2002 Turing Award lecture, "Cryptography is typically bypassed, not penetrated." I would argue that this is true of all aspects of cybersecurity mechanism including those delivered by building security in; it is the possibility of bypass that ultimately matters. Our sentient opponents know that, too, and their investments in automating the discovery of methods of bypass are in a hell of a horse race with both building security in and in static analysis of code bodies, new or old.

I pause here to add that colleagues in the forefront of static analysis report that they are seeing web applications in excess of 2GB with 20K variables, applications that can only have been written by machine, yet they, too, have flaws. Perhaps if the building in of security is to remain an embodiment of a paradigm of defense, there will have to soon be research puzzles addressing how to prevent machines from writing vulnerabilities.

As with other paradigms in crisis, the crisis would be vacuous if there were not at least the possibility of a paradigm to replace the one that is creating and thus confronting anomalies. A specific one that I feel holds promise is the work going on at U Penn by Clark, et al., on what they call "the honeymoon effect." [SC] In rough terms, they quantify the degree to which sentient opponents require non-zero time to exploit new code and, in turn, how to simply outrun those opponents, and in so doing they also quantify how code re-use reduces the labor of exploit. Is that an alternative paradigm to defense, at least to the embodiment of defense in the form of hardening deployed systems? Time will tell, but I believe that an alternative paradigm could be in the running though, as we all recognize, constant code churn is fundamentally inconsistent with compliance and certification.

The paradigm of defense had a similar challenge some years ago under a different regime of tight coding, namely the transition whereby Microsoft, in particular, adopted address space layout randomization (ASLR) to thwart, in particular, the problem of buffer overflows. Perhaps moving target defense should be classified in a similar way, namely a new solution to the puzzle of how to prevent exploitability. I think it goes beyond that, but that is an arguable nuance at least for now. But both randomization of code bodies at run time and moving target by way of rapid release embody Einstein's wisdom encapsulated in his remark that "Insanity [is] doing the same thing over and over again and expecting different results." Where we are losing, we have to change the rules of the game. Just bearing down is to make adding epicycles to Ptolemaic equations the puzzles we are solving.

Speaking from my engineering bias, per se, for me the pinnacle goal of cybersecurity engineering is that of "No silent failure." Failure is to be avoided whenever possible, but absolute total avoidance is sure to be diseconomic and not of interest in the deployed world. Silent failure is pandemic. The Verizon Data Breach Investigations Report [DBIR] has time and again found that 80% of all data breaches are discovered not by the victim but by an unrelated third party. With a colleague, we run the Index of Cyber Security [ICS] and in that we once asked whether the respondent had ever discovered a data breach of another firm not his/her own. We got 55% "yes and confirmed" and 10% "yes but unconfirmed" for a total of 65%. Since we exclude law enforcement from the Index's catchment, our 65% and

Verizon's 80% can be said to be in close agreement -- there is a lot of silent failure of data protection out there. As you well know, the lay press regularly reports silent failures of data protection in the retail sector, but is it not likely that the more severe the data protection failure the more likely it is to be silent, that is to say to \*still\* be silent?

I was taught a long time ago, but the rule of thumb I learned then was that for a large code base a substantial portion (I was taught 40%) of it should be in exception handling. I now interpret thoroughgoing attention to exception handling as the avoidance of silent failure. Perhaps that is the change of paradigm that is needed, to ensure not that code cannot and will not fail, but that failure will not be silent. Certainly the language theoretic security work going on at Dartmouth[LANGSEC] and elsewhere has a similar view and it, too, has promise if indeed the crisis in code security is severe enough to require a new paradigm consistent with the idea of no silent failure. The core sentence in their manifesto is that "...the only path to trustworthy software that takes untrusted inputs is treating all valid or expected inputs as a formal language, and the respective input-handling routines as a recognizer for that language. The recognition must be feasible, and the recognizer must match the language in required computation power." That is what an alternate paradigm sounds like.

Kuhn takes some pains to say why it is that a paradigm shift requires a crisis, that is that "to an extent unparalleled in other fields, [scientists] have undergone similar educations and professional initiations." One here must ask the central question of this essay by mirroring Kuhn, are the paradigms of cybersecurity in enough of a crisis that resolution of the crisis requires a change of paradigm? The answer is by no means obvious, though to my eye there is a crisis or, rather, several crises now in play. If the crisis is, or the crises are, sufficient to require a reformulation of the paradigm or paradigms of cybersecurity, then a scientific revolution is upon us, what Kuhn calls "a reconstruction of group commitments." As he points out, a crisis requiring such a reconstruction may not even be in cybersecurity itself, but instead due to discoveries in some other field or venue, just as discoveries in physics engendered a crisis in chemistry once upon a time.

On the other hand, perhaps I am being too hard on cybersecurity. Perhaps it is already more mature than I give it credit for. In Kuhn's analysis, a mature science has (1) a relative scarcity of competing schools, (2) members of the scientific community provide the only audience for, and judges of, what constitutes puzzles worth solving, and (3) puzzle solving is the principal activity in which the scientists are engaged. In such a situation, "scientific knowledge is intrinsically the common property of [the] group or else nothing at all." Certainly not everyone agrees with that characterization, not even with the idea of puzzle solving -- Prof. Peter Drucker famously said "Don't solve problems. Create opportunities." While that is not research, might we consider a paradigm crisis to be the kind of problem for which the formulation of a new paradigm creates opportunities? Or take Francis Bacon, "Truth emerges more readily from error than from confusion." Are we at a point of either error or confusion and, if so, has the reconstruction of what we are about become timely?

Kuhn's analysis was, of necessity, grounded in the history of several centuries of science. Perhaps he agreed with Winston Churchill,

that "The further back we look, the further forward we can see." But the centuries that both Kuhn and Churchill absorbed may be misleading in a world of accelerating change. Kuhn maintained over and over that science is not linear, not some steady upslope at 8% grade. Gould said the same thing about biologic evolution. One can only assume that the science of security will inevitably experience periods of relative stasis, what Kuhn called normal science, equilibria punctuated by periods of rapid change. But what if rapid change is a constant? What if the periods of expansion and consolidation of what a paradigm allows us to scientifically solve become themselves too short for thoroughly exploring what our then current paradigm empowers us to do? "Normal science does not aim for novelties of fact or theory and, when successful, finds none."

Let me quote a longer passage that challenges us particularly. "In the development of any science, the first received paradigm is usually felt to account quite successfully for most of the observations and experiments easily accessible to that science's practitioners. Further development, therefore, ordinarily calls for the construction of elaborate equipment, the development of an esoteric vocabulary and skills, and a refinement of concepts that increasingly lessens their resemblance to their usual common-sense prototypes. That professionalization leads, on the one hand, to an immense restriction of the scientist's vision and to a considerable resistance to paradigm change. The science has become increasingly rigid. On the other hand, within those areas to which the paradigm directs the attention of the group, normal science leads to a detail of information and to a precision of the observation-theory match that could be achieved in no other way. Furthermore, that detail and precision-of-match have a value that transcends their not always very high intrinsic interest. Without the special apparatus that is constructed mainly for anticipated functions, the results that lead ultimately to novelty could not occur. And even when the apparatus exists, novelty ordinarily emerges only for the [scientist] who, knowing with precision what he should expect, is able to recognize that something has gone wrong. Anomaly appears only against the background provided by the paradigm. The more precise and far-reaching that paradigm is, the more sensitive an indicator it provides of anomaly and hence of an occasion for paradigm change. In the normal mode of discovery, even resistance to change has a use... By ensuring that the paradigm will not be too easily surrendered, resistance guarantees that scientists will not be lightly distracted and that the anomalies that lead to paradigm change will penetrate existing knowledge to the core. The very fact that a significant scientific novelty so often emerges simultaneously from several laboratories is an index both to the strongly traditional nature of normal science and to the completeness with which that traditional pursuit prepares the way for its own change."

That, then, is the question before us, complicated by the changing nature of what scientists of security are studying both with respect to rapid technologic change and the presence of sentient opponents, leavened, of course, with the societal demands fast upon us largely independent of what we know or say. I think I see paradigms here that are, or soon will be, in undeniable crisis. I can, of course, be entirely wrong and we may still be working our way up to being a science, still coalescing schools of thought into the kind of paradigm that will define us as scientists.

Perhaps some one of you or your colleagues is already further down the track of applying Kuhn or something like it to the science of security. If so, I am eager to hear all about it. That I do not know about it despite it being already well underway is merely my error, and as quoted above, perhaps one of you will teach me what that error has made me receptive for.

It is hard to tell that you are in the knee of the curve, so I will close with a line from the musical Gigi, "Have I been standing up too close or back too far?"

There is never enough time. Thank you for yours.

=====

[TSK] T.S. Kuhn, *The Structure of Scientific Revolutions*, 1962 & 1969

[JB] John C. Bogle, *The Clash of the Cultures: Investment vs. Speculation*, 2012

[NNT] Nassim Nicholas Taleb, *Fooled by Randomness*, 2004

[SJG] Stephen Jay Gould, *Wonderful Life*, 1989

[MS] Matthew R. Stoneking, lecture, Lawrence Univ., 2001

[BG] Brian Greene, *The Elegant Universe*, 2000

[SOS] NSA, Science of Security  
[www.nsa.gov/public\\_info/press\\_room/2014/Best\\_Scientific\\_Cybersecurity\\_Paper\\_Competition.shtml](http://www.nsa.gov/public_info/press_room/2014/Best_Scientific_Cybersecurity_Paper_Competition.shtml)

[CD] letter to Henry Fawcett, 1863, as quoted in *Medicine in Quotations: Views of Health and Disease Through the Ages*, Huth & Murray, ed., 2006

[DG] <self>, "The Shrinking Security Perimeter," audio mirror at [geer.tinho.net/Dan\\_Geer\\_-\\_The\\_Shrinking\\_Security\\_Perimeter.mp3](http://geer.tinho.net/Dan_Geer_-_The_Shrinking_Security_Perimeter.mp3), 2004

[EH] Eric Hughes, "A Cypherpunk's Manifesto," 1993

[DG2] <self>, "Tradeoffs in Cyber Security," UNCC, 2013  
[geer.tinho.net/geer.uncc.9x13.txt](http://geer.tinho.net/geer.uncc.9x13.txt)

[DS] Daniel Solove, *Nothing to Hide*, 2011

[SFI] "Optimality vs. Fragility: Are Optimality and Efficiency the Enemies of Robustness and Resilience?"  
[www.santafe.edu/gevent/detail/business-network/1665](http://www.santafe.edu/gevent/detail/business-network/1665)

[DARPA] Cyber Grand Challenge,  
[www.darpa.mil/Our\\_Work/I20/Programs/Cyber\\_Grand\\_Challenge\\_\(CGC\).aspx](http://www.darpa.mil/Our_Work/I20/Programs/Cyber_Grand_Challenge_(CGC).aspx)

[SC] Sandy Clark, Michael Collis, Matt Blaze, Jonathan Smith, "Moving Target: Security and Rapid-Release in Firefox," 2014,  
[dl.acm.org/citation.cfm?id=2660320](http://dl.acm.org/citation.cfm?id=2660320)

[DBIR] Verizon DBIR, [www.verizonenterprise.com/DBIR](http://www.verizonenterprise.com/DBIR)

[ICS] Index of Cyber Security, [cybersecurityindex.org](http://cybersecurityindex.org)

[LANGSEC] [www.langsec.org](http://www.langsec.org), consider attending the workshop in San Jose to be held on Thursday May 21, 2015

=====

See also

American Chemical Society, "Thomas Kuhn Paradigm Shift Award"  
[www.materialsviews.com/2013-thomas-kuhn-paradigm-shift-award](http://www.materialsviews.com/2013-thomas-kuhn-paradigm-shift-award)

And consider the paradigmatic issues around long, long periods of overlap between fielded systems built under differing paradigms, such as

- . versions of SSL/TLS
- . magstripe -> EMV payment
- . IPv4 -> IPv6 networking

=====

This and other material on file at <http://geer.tinho.net/pubs>