Extracting Unlearned Lessons from Past Poor Choices lest They be Learned the Hard Way in the Future

Dan Geer

Unlearned lessons are those insights missed from a past situation. When we do not learn from experiences, we continue to make the same decisions in similar situations. In the case of the United States, unlearned lessons undermine the future security and prosperity of democracies. The results of unlearned lessons can be the individual's free choice, but others, including some facing us now, heavily burden the future with the collective history of other prior choices. More volatile times face open societies globally. As Nassim Taleb observes, when the tails of a probability distribution get fatter, the predictable becomes a function of the distribution's extreme values and only those extreme values.[1] In multiple publications, Taleb argues that the world is "undergoing a switch between continuous low-grade volatility to a process moving by jumps, with less and less variations outside of jumps."[2] The faster the rate of systems change, the heavier become those tails, due mostly to the growth of unrecognized interdependence between the moving parts. In the statistical analysis of systems, if one is uncertain about the tails of the data, then one is uncertain about the mean as well. Yet, the faster the rate of systems change, the heavier become those tails, due mostly to the growth of unrecognized interdependence between the moving parts, and thus the less useful for learning are their means. Such a situation requires the prudent person to plan for maximal damage scenarios, not for most probable scenarios, and to ensure that the choices they make along the way offer reasonable and secure alternatives when the worst scenarios emerge.

Rather than exploring one or two broad unlearned lessons, as in other articles in this publication, this essay presents a series of questions in an arbitrary order whose answers could lead to different possible futures. These questions, and the choices made

All ideas stated here are solely those of the author(s) and do not reflect the positions or policies of any element of the U.S. Government.
© 2021 Dan Geer

Dan Geer, Senior Fellow, In-Q-Tel. Milestones: The X Window System and Kerberos (1988), the first information security consulting firm on Wall Street (1992), convenor of the first academic conference on mobile computing (1993), convenor of the first academic conference on electronic commerce (1995), the "Risk Management is Where the Money Is" speech that changed the focus of security (1998), the Presidency of USENIX Association (2000), the first call for the eclipse of authentication by accountability (2002), principal author of "Cyberinsecurity: The Cost of Monopoly" (2003), co-founder of SecurityMetrics.Org (2004), convener of MetriCon (2006-present), author of "Economics & Strategies of Data Security" (2008), and author of "Cybersecurity & National Policy" (2010). Creator of the Index of Cyber Security (2011) and the Cyber Security Decision Market (2012). Lifetime Achievement Award, USENIX Association, (2011). Expert for NSA Science of Security award (2013-present). Cybersecurity Hall of Fame (2016) and ISSA Hall of Fame (2019). Testified five times before Congress.

in answering them, are not commonly addressed in the broader literature. This essay intends to capture choices in national cybersecurity of which we seem unaware, yet which have great steering power that, ten years from now, will seem obvious in retrospect. By then, some of the choices that follow below will have already been made, choices that are expensive to reverse in either dollars or clock-ticks. These questions are neither mutually exclusive nor collectively exhaustive, but each marks a fork in the road and a lesson yet to be learned about the complexity of our society, choices, and futures. Each question is chosen to illustrate an extant fork in the road to the future.

SURVEYING UNLEARNED LESSONS IN UNAPPRECIATED PAST FORKS IN THE ROAD

Automation: Over time, does automation help defense or offense more?

If the answer over time is offense, then cybersecurity becomes more like nuclear deterrence, including that there are few precision weapons and that enforceable treaties are essential to our survival. If the answer is defense, then we should turn to algorithms to do what we cannot otherwise do, that is, to protect us from other algorithms.[3]

Connectivity without Management: How much do we allow self-modifying decision-making to free-run?

The Berkeley SWARM lab says that there will be 1,000 radios per person on earth within five years, while Pete Diamandis (X Prize) says within 15 years, there will be 10,000 sensors per person. That scale exceeds manageability, so that technology must just be allowed to free-run. Yet, if that free-running is self-modifying, then it will be unable to be evaluated for trustworthiness. The choice is to what extent we will pair free-running with self-modifying cybersecurity systems, and the challenge then will be by what mechanisms will we make that choice stick.

Metrics as Policy Driver: Do we steer by Mean Time Between Failure (MTBF) or Mean Time To Repair (MTTR) in Cybersecurity?[4]

Choosing Mean Time Between Failure (MTBF) as the core driver of cybersecurity assumes that vulnerabilities are sparse, not dense. If they are sparse, then the treasure spent finding them is well-spent so long as we are not deploying new vulnerabilities faster than we are eliminating old ones. If they are dense, then any treasure spent finding them is more than wasted; it is disinformation. Suppose we cannot answer whether vulnerabilities are sparse or dense. In that case, a Mean Time To Repair (MTTR) of zero (instant recovery) is more consistent with planning for maximal damage scenarios. The lesson under these circumstances is that the paramount security engineering design goal becomes no silent failure – not no failure but no *silent* failure – one cannot mitigate what one does not recognize is happening.

Embedded Systems: Do we require devices to either be updatable or have a finite lifetime?

A remote management interface is required for updatability, but such an interface is itself a potential risk should it be cracked.[5] Yet, doing without remote updates risks having devices that are immortal and unfixable unless a finite lifetime is purposefully designed in. Enforcement of either mandate—that of a remotely fixable or pre-limited lifetime—is an expressly government issue, and enforcement means culpability for those who fail to implement the one or the other. But the choice then becomes hinged on how that culpability is defined and enforced. One lesson to note is that devices that are unlocatable after deployment are a special case where remote updates are not possible; these devices must have limited lifetimes.

Composability: Since non-composability—where secure A plus secure B does not make a secure A+B—is common, how is research directed?

Composability is both a research and commercial opportunity that allows cybersecurity to aggregate as the sum of secure parts. This is the inherent security-in-the-limit question for the supply chain model. Within the narrow domain of just input parsers, for example, hostile data input is so very often the mechanism by which exploit is triggered.[6] Composable parsing would yield much general benefit, and not just for cybersecurity if research were to resolve the mechanisms of its creation or at least the detection of cybersecurity non-composability in critical supply chains. With some encouragement, DARPA has become interested in looking at composability, but more attention is needed.

Reproducible Buildability: Do we mandate that the software supply chain be checkable, or should the buyer absorb all the risk?

Olav Lysne has shown that, in the limit, electronic equipment from untrusted vendors cannot be verified.[7] If we choose not to let the buyer absorb all the risk, then the fallback choice is to introduce process requirements, requiring merchants of digital goods to retain the ability to recreate those digital goods exactly as deployed. This obligation requires that

the build environment is preserved in working detail. A nascent, international effort on "reproducible builds" is making technical progress but more support is needed to make this choice a reality.[8]

Abandonment: Every kind of property save software can be seized in the public interest if and when abandoned; how do we define and remedy software abandonment going forward?

When a supplier declares that a software product is "no longer supported," the software should be considered abandoned, thus triggering some legal seizure or collection process. The most straightforward remedy would be to open-source abandoned codebases. Another would be to seize cryptographic keys of signal value, such as those of a certifying authority that goes out of business. The nascent Microsoft/GitHub arctic repository of code on which "we the people" depend is an example answer of serious preservation, but it is not a legal answer applicable across national systems. The lessons of abandoned and unsupported software in the wild already are well demonstrated, making this question another in which choices here channel the future and have yet to be learned.[9]

Attribution: Do we geocode the Internet by fiat or keep arguing about attribution?

Nationally, a public safety argument drove the requirement to geocode mobile phones in real-time. The sanctity of all financial services already starts with "know your customer," including from where they normally connect to the Internet. Whereas intercontinental ballistic missiles have a visible flight path and a limited number of launch-capable governments, yet offensive software has neither. Would geocoding the Internet offer a simple attribution resolution? The Westphalian principle is clear for state responsibility; if packets are allowed to exit a nation's territory, then that country must have strict liability for them regardless of any actual negligence or intent to harm on that nation's part. Despite concerns with adversaries' engaging in false flag campaigns by routing traffic through other nations, the greater harm lies in not choosing to at least know the origins of traffic entering and exiting nations.[10]

Prioritization: Does continuous machine learning require prioritizing integrity over confidentiality?

Confidentiality has long been first among equals within the classic triad of confidentiality, integrity, and availability (CIA). As data collection grows in breadth and depth, and as data-driven algorithms take broad control of the physical world with the corruption of inputs among the larger threats, must not data integrity take priority over availability and availability take priority over confidentiality? If this is the moment of overturning that prior order, then much different design regimes must be deployed, including, but not limited to, robust data chain of custody regimes uninfluenceable by single points of failure.[11]

Data Mitigation: What mitigation, if any, can be enforced for data loss immortalized in immutable storage (such as blockchains)?

Inherent to blockchain design, the insertion of bad things onto blockchains has no reversal mechanism, no "undo" is available. Do those who manage blockchains without a deletion mechanism inherit strict liability for the consequences of the blockchain's misuse as the publisher of illicit, inaccurate, or harmful goods? If they do not, then who is liable and for how long?[12] Do democratic societies proscribe certain information from being put in immutable storage?

Namelessness: Do we require resolvable names for some classes of services?

At present, name resolution is the double-entry bookkeeping of the internet. In this process, if A claims this specific name and URL address pair are connected, then B can verify A's claim with third parties. In the absence of verifiable names, however, how does anyone verify to whom they are connected? If names are not required, a wide range of subordinate questions whose answers have considerable impact emerge because one side of the verifying bookkeeping is now missing. Should a (nameless) address be permitted to be baked into code and, if so, where? Are software updating engines supposed to reach out to devices that have only a network address and no name? What happens if an endpoint that is asked to accept an update has no name to check? Do we need to develop a two-class liability system and/or duty to support name verification, one for checkable name-to-address pairs versus one for un-named addresses? Will internal firewalls now have to include a key-centric, rather than a name-centric, PKI? Does a MAC address or a UUID-in-ROM distinguish keys in a nameless world and thus offer an alternative to an identity-based PKI tied to hardware? Either way, these choices have implications for a range of internet security management, especially as to whether the key-management job is going to be harder or easier absent names.[13] What would be the risk interaction of a nameless internet with an IPv6 world too big to enumerate, especially when IPv6 protocols inherently allow address hopping and multi-homing, challenging existing security regimes?

Interrogability: Does one accept dependence on black-box algorithms, and if so, to what extent?

Machine learning from noisy datasets delivers equations where coefficient values have no inferential meaning and are likely to have the fragility that comes from over-fitting. Algorithmic decision-making is being increasingly challenged by new rules and regulations around the world that are attempting to increase transparency and limit disfavored outcomes. Article 15 (and others) in the European Union's General Data Protection Regulation (GDPR), for example, creates a right to challenge any algorithmic decision by asking "why?", which cannot be answered by an uninterrogable algorithm.[14] Does this create a new variety of "informed

consent" when a medical algorithm says what must be done but cannot explain why? And so what, exactly, does "informed" in "informed consent" then mean? Similarly, military users will likely be reluctant to permit algorithms to apply deadly force if those algorithms cannot explain themselves in the vernacular. At what level of criticality do we require interrogability prior to legal deployment?

Insurance: What is the definition of force majeure in cybersecurity?

As Mondelez v Zurich and Merck v Allianz/AIG show, what is force majeure in cybersecurity may be judiciable. Since many instances of attack have no immediately identifiable attacker, how will this work out? Has the burden of proof for payout shifted from proving "I was harmed" to proving "I was harmed, and I can prove that I was the actual target"? Besides that, how will risk pooling—the very premise of insurance—proceed in the future? At some level of personalization, there is nothing left with which to pool.[15]

Entity scale: Are cloud computing suppliers critical infrastructure?

The very smallest companies and the very biggest companies are safer if they do their computing on-premises, whereas middling companies are safer if they do their computing in the cloud.[16] In a target-rich world, the smaller organizations do not draw as many attacks, while the bigger ones have the dollars to invest in stronger cybersecurity and corner the talent market. The middle-sized organizations are, therefore, all but obliged to have somebody else mind the cookstove. So, does it then follow that cloud computing suppliers have the duties of a critical infrastructure provider? Do we replicate the regulatory structure surrounding Systemically Important Financial Institutions (SIFI) for computing entities that are not "too big to fail" but rather "too interconnected to fail"? What is the policy analog for a cloud provider to SIFI's capital adequacy requirements or any of the rest of the financial stress-testing regime? Since the bulk of regulation for financial services is the accumulation of past failures, shall we just wait for cloud failures of sufficiently serious kinds to occur and then design appropriate regulations?

Failure as design: When a functionality is a security failure by-design, who pays the price?

Repeatedly, we feign surprise when an ICT product or service is found to have design failures. As an engineer would say, what might be possible and consistent with the paradigm of risk tolerance is that no system may fail silently. In a sense, that is what data breach laws assume—that breaches will happen no matter what and that the proper response is remediation and notification of affected parties. If users consider notification adequate mitigation, then this type of law does create a form of security as the surprise is followed by its mitigation. Any regulation in this area would include performance standards for a latency of cleanup steps, like notification.

Moreover, as a Verizon Data Breach Report noted, 3/4 of all data losses are discovered by unrelated third parties, and whether data breaches are preponderantly insider attacks or outsider attacks depends on the definition of an "insider." If "insider" means "on the payroll," then insider attacks are not the most important issue. If, however, insiders include folks on the payroll, plus employees of the organization's partners or vendors who have as-of-right access to that company's infrastructure, plus employees of those in its supply chain, then the majority of data losses are insider attacks.[17]

Ownership: Should the beneficial owner of software also own its risks?

Today, suppliers claim ownership for the software they sell and license to users, yet via the pervasively unreadable End User License Agreement, they disclaim all responsibility for what they say they own. This contradiction now applies to countless mélanges of hardware and software from the biggest to the smallest products. This can either continue toward its logical conclusion, that users own nothing and have no recourse, or it can be steered (at least in democratic countries) toward some form of software liability. To do that, new constructions of strict liability and merchantability are necessary, especially considering self-modifying software where the copy a user has may well be unique.[18] This fork in the road demonstrates, more than others, that does not make a clear decision in this instance is equivalent to having made a decision. A plague of lobbyists can be expected to join the fray here.

Truth: What are the trust anchors we now need?

While the common use of the term "trust anchor" refers to the self-signed certificate at the head of a specific cryptographic hierarchy, here it is intended in a policy and societal sense. As Malcolm Gladwell has persuasively argued, civil society only works if its members can safely default to trust in dealing with each other.[19] As the digitalization of society proceeds, the presence of cyber risk makes defaulting to trust look increasingly naive. What divergence effects do you get if, among humans, you more or less have to default to trust, whereas in bitspace, you more or less have to default to mistrust or, as it is now commonly referred to, adopt a "zero trust" policy?

Sovereignty: Is it more prudent for the US to ban certain technology (e.g., Huawei) or to do what it takes to have technical leadership?

The intersection of patents and standards, and the coercion or protection applied to intellectual property distinguishes an authoritarian system versus a democratic system. The global suppliers will not take sides out of unpressured patriotism. For small countries, the choice is mostly whose technology to adopt and thus what side-effects to endure. For great cyber powers, the choice is whether to do what it takes to win the technology race. Furthermore, democratic nations do things with allies, not to them.

Balkanization: Should the trade in network and security gear mimic the trade in arms, that is to say, that you only buy or sell within your allies or region?

This question concerns whether all sellers become holding companies so that, for example, the Apple-China team does one thing while the Apple-EU team does another. A government official saying "your government needs you to do us a favor" is probably a seller issue more than a buyer issue. Even if there is no nation-state pressure on a supplier to install an exploitable flaw, that "do us a favor" can take the form "you are forbidden to fix the following flaw." Or is nation-state the wrong granularity here? Some Google security engineers say they only trust Google-manufactured phones as other manufacturers modify the Android software and may install other software as part of the base install. A few major banks no longer buy application software, they write their own. As cyber sovereignty spreads, it could come to require that one runs only code that one or one's allies directly wrote.

Strategy: What strategy is most cost-effective for cybersecurity: cost-efficiency or cost-effectiveness?

Cost/benefit analysis is useless and wasteful for cybersecurity; what matters is cost-effectiveness. The fork in this road is to prove that a body of code implements its specifications and no more than its specifications or to embrace a moving target defense. Code proofing, the first option, is a—if not the—gold standard, but it is also difficult to do. After a successful proof is in hand, an organization must adopt strict change control to protect its investment. The second option is moving target defense, which is comparatively easy and may work well enough to diminish today's strategic asymmetry that favors attackers over defenders. However, it also guarantees that how any fielded instance of code actually works on any endpoint is not obvious. It requires an intermediary to explain the process. Brutal change control – the first option – damps down the rate of change, which may be an acceptable price to pay for stamping out cyber exploits and other vulnerabilities. Moving target defense does not care if there is a new set of vulnerabilities with every revision, which may be an acceptable price to pay for stamping out cyber exploits and other vulnerabilities.[20]

Taking either path is a choice supported by results obtained by serious scientists; it is also one of those choices that are expensive in either dollars or clock-ticks to later reverse.

Analog fallback: If a state of security is the absence of unmitigable surprise, what is the mitigation for an irrecoverable software fault?

Of all the forks or choices enumerated in this article, this one is the most telling choice. Optimality and efficiency work counter to robustness and resilience. Complexity hides interdependence, and unacknowledged interdependence is the source of black swan events. The benefits of digitalization are not transitive, but the risks are. Because single points of failure require militarization wherever they underlie gross societal dependencies, frank minimization

of the number of such single points of failure is a national security obligation. Cascade failure ignited by random faults is quenched by redundancy, whereas cascade failure ignited by sentient opponents is exacerbated by redundancy. Therefore, the preservation of uncorrelated operational mechanisms is likewise a national security obligation. Those uncorrelated operational mechanisms are the analog alternative.

The ability to operate in the absence of the Internet and all that it delivers requires retaining the pre-digital analog realm is crucial. It does not share common-mode failures with the digital realm, nor does analog create pathways to cascade failure. To retain the analog realm, it must be updated, used, and exercised, and not left sitting on some societal shelf hoping that it still works when some emergency demands that it work. As a current example, consider the regulatory and public policy debate ongoing in Sweden on how to retain a sufficient baseload of cash processing such that the mechanisms needed to use cash do not disappear, even if society is moving toward a cashless system.[21] Sweden is now choosing whether to make cash processing a critical infrastructure and to do so by regulation. Mere economics will otherwise create a singleton risk of society-wide failure.

In short, the intervention is to require, by force of law, that those who choose to opt-out of the digital world do not do so at the price of moving back to the 15th century, that accommodating those who opt-out is precisely and uniquely how to maintain a baseload for nearly all the essential components of our daily living. Perhaps never before and never again will national security and individual freedom jointly share a call for the same initiative at the same time. The active preservation of the analog option must come soon if it is to have the cost-effectiveness of preserving fully amortized infrastructure and not the sky-high costs of recreating it under emergency conditions.

THE UNDERLYING UNLEARNED LESSON: KEEP ANALOG OPTION HEALTHY AND AVAILABLE

As mentioned at the outset of this essay, the prudent person should plan for maximal damage scenarios, not for the most probable scenarios, and avoid creating feelings of safety that are unwarranted.[22] Most important ideas are not exciting. Most exciting ideas are not important. Not every problem has a good solution. Every solution has side effects. To learn the unlearned lessons of the past, it is best to have a prudence doctrine and anticipate possible future losses but not future gains.[23]

While change is inevitable, undirected change does not produce some steady upslope at 8% grade; instead, the course of undirected change naturally traces through periods of quietude interrupted by bursts of self-reinforcing change. The evolutionary biologist Stephen Jay Gould called this phenomenon a “punctuated equilibrium.”[24] By this author’s count, there have been three “equilibrium punctuations” to date. The first was circa 1995: a TCP/IP stack

appeared as a freebie in Microsoft Windows, exposing an unhardened operating system to the world, a change that birthed the cybersecurity industry. Second, circa 2005: cybered opponents changed from braggarts to professionals, a critical change because while braggarts may share their tools, professionals hold them closely. The professional's ever-increasing stockpile of software exploits is why zero-days have come to dominate risk. Third, circa 2015: blind automation of flaw finding—which was a DARPA Cyber Grand Challenge intended to help defense—is becoming a central tool in the attackers' kit and is perhaps the ultimate proof that all security technology is dual-use. Looming is a fourth punctuation, possibly circa 2025, as a wide variety of developments in emerging technology and geopolitics converge. The choices being made now will be as consequential as any that have been made before.

At the end of the day, ordered liberty depends on putting a speed limit on irrevocable change. We cannot be passive.

NOTES

1. Pasquale Cirillo and Nassim Nicholas Taleb, "What are the chances of a third world war?" *Significance* 13, no. 2 (2015).
2. Nassim Nicholas Taleb Pasquale Cirillo, "On the Statistical Properties and Tail Risk of Violent Conflicts," *Physica A: Statistical Mechanics and its Applications,* no. 429 (2016), https://arxiv.org/abs/1505.04722.
3. Https://spw18.langsec.org/slides/Vanegue-AEGC-5-year-perspective.pdf.
4. Https:// www.theatlantic.com/technology/archive/2014/05/should-hackers-fix-cybersecurity-holes-or-exploit-them/371197.
5. https:// tools.ietf.org/html/rfc791.
6. https:// langsec.org.
7. Olav Lysne, *The Huawei and Snowden Questions: Can Electronic Equipment from Untrusted Vendors be Verified? Can an Untrusted Vendor Build Trust Into Electronic Equipment?* (Springer Nature, 2018).
8. https:// reproducible-builds.org.
9. http:// geer.tinho.net/ieee/ieee.sp.geer.1307.pdf.
10. http:// geer.tinho.net/ieee/ieee.sp.geer.1707.pdf.
11. http:// geer.tinho.net/ieee/ieee.sp.geer.1801.pdf.
12. https:// news.bitcoin.com/a-brief-history-of-hidden-messages-in-the-bitcoin-blockchain.
13. http:// geer.tinho.net/fgm/fgm.geer.1812.pdf.
14. https:// www.clarip.com/data-privacy/gdpr-article-15.
15. www.bloomberg.com/news/features/2019-12-03/merck-cyberattack- s-1-3-billion-question-was-it-an-act-of-war
16. geer.tinho.net/fgm/fgm.geer.1909.pdf; and geer.tinho.net/fgm/fgm.geer.1412.pdf.
17. http://geer.tinho.net/ieee/ieee.sp.geer.1911.pdf, http://geer.tinho.net/fgm/fgm.geer.2006.pdf; http://geer.tinho.net/fgm/fgm.geer.2012.pdf; http:// geer.tinho.net/fgm/fgm.geer.1906.pdf.
18. The End of Ownership, Perzanowski & Schultz, MIT, 2019, and http://geer.tinho.net/ieee/ieee.sp.geer.1907.pdf.
19. Talking to Strangers, Gladwell, Hachette Book Group, 2019.
20. https://www.ll.mit.edu/sites/default/files/page/doc/2018-05/22_1_8_Okhravi.pdf.
21. https://www.ft.com/content/c4ea6fb3-6262-4426-9503-0539lf0e523a.
22. http://geer.tinho.net/geer.mindthesec.16ix20.txt.
23. https://www.hoover.org/research/rubicon.
24. Niles Eldredge & S.J. Gould, "Punctuated equilibria: an alternative to phyletic gradualism" In T.J.M. Schopf, ed., Models in Paleobiology. San Francisco: Freeman Cooper, 82-115 (1972); Stephen Jay Gould and Stephen Jay Gould, *Punctuated equilibrium* (Harvard University Press, 2009).