

A Rubicon

DANIEL E. GEER, JR.

Aegis Series Paper No. 1801

We state as an axiom that cybersecurity and the future of humanity are now conjoined, and that that conjoining makes the scope of a full analysis of the situation too broad and too subtle for an essay as short as this one.

Time is irreversible. To make something timeless is to make it irreversible. We want our security assertions to be irreversible meaning timeless, or timeless meaning irreversible. We invent security mechanisms to get irreversibility and, by way of that, timelessness. We calibrate our inventions' quality by the imagined time that must elapse before they can be undone, which is to say reversed. The inventions that deliver "You are safe here for a millisecond" and "You are safe here for a millenium" differ in the time constants of their reversibility.

Such time constants impose an ordering on alternate inventions, the longer the better, all other things being equal. Each and severally, the curves of implementation cost and the notional value of timelessness cross somewhere, a somewhere we will discuss.

Because the wellspring of risk is dependence, aggregate risk is a monotonically increasing function of aggregate dependence. We call that on which we most depend "critical infrastructure(s)." Because dependence is transitive, so is risk. That you may not yourself depend on something directly does not mean that you do not depend on it indirectly. We call the transitive reach of dependence "interdependence," which is to say, correlated risk.

Interdependence within society today is centered on the Internet beyond all other dependencies excepting climate, and the Internet has a time constant of change five orders of magnitude smaller than that of climate. Our concern is unacknowledged correlated risk, the unacknowledged correlated risk of cyberspace is why cyberspace is capable of black swan behavior. Unacknowledged correlations contribute, by definition, to heavy tails in the probability distribution of possible events. "For fat-tailed variables, the mean is almost entirely determined by extremes. If you are uncertain about the tails, then you are uncertain about the mean."¹ As complexity hides interdependence(s), ergo complexity is the enemy of security. (For precision, "complexity characterises the behaviour of a system or model whose components



interact in multiple ways and follow local rules, meaning there is no reasonable higher instruction to define the various possible interactions.”)²

There are two—precisely and only two—classes of cyber risks that directly rise to the level of national security. One class is those critical services which by the very definition of their mission must create a single point of failure. (The integrity of the Domain Name system root is an example.) For design-implicit single points of failure, defense in depth is our only answer. It is not a research grade problem—it is a referendum on the state’s willingness to spend money.

The other risk at the level of national security is that of cascade failure which implicates services that are critical or merely pervasive. Cascade failure is where complexity and interdependence intersect, where unacknowledged correlation lies, where the heavy tails await.³ We already know that optimality and efficiency work counter to robustness and resilience and vice versa.⁴ We already know that a “control” strategy fails worse than a “resilience” strategy. Quenching cascade failure, like quenching a forest fire, requires an otherwise uninvolved area to be cleared of the mechanisms of transit, which is to say it requires the opposite of interdependence. Senate Bill 79 addresses the national security priority of operating the most critical infrastructures by analog, disconnected means.⁵ As such, S.79 is precisely the kind of step that must be taken. In the absence of purposeful disconnectedness at critical points, the mechanics of complexity then take hold so that we undergo “a switch between [continuous low grade volatility] to . . . the process moving by jumps, with less and less variations outside of jumps.”⁶

It is conservative to assume that the hardest lessons are learned the hardest of ways. But let us hope we can, at least, muster the resolve to simulate what failures we can imagine.⁷

If a dependence creates a risk, then logic tells us to either forego the dependence or mitigate the risk. Per the present author’s definition, a state of security is the absence of unmitigatable surprise—there will always be surprises, but the heavy tails that accompany complexity mean that while most days will be better and better, some days will be worse than ever before seen. The provision of mitigation thus becomes more challenging over time as complexity accumulates and unacknowledged correlated risks become embedded.

Mitigation is always context dependent. In settings where failures come from the slow degradation of mechanical systems, redundancy is the straightforward mitigation

technique. Where failures come from intentional actions by sentient opponents, redundancy adds to risk rather than subtracting from it because redundancy replicates vulnerability to sentient action but not to random events. To thwart the outcomes that follow sentient opponent actions, diversity of mechanism is required. Cascade failure is not quenched by redundancy but by the provision of required function via alternate mechanisms that do not share common modes of failure, that is to say mechanisms that do not interdepend and thus do not contribute heavy tails to the distribution of failure probabilities.⁸

Much of our challenge follows from the rapid rate of change which we otherwise praise—Innovation Uber Alles, All Hail. Documented a decade ago by Ozment and Schechter, we know that stable code bases under stable oversight can cleanse themselves of vulnerabilities over time.⁹ Clark et al. have since shown in measurable ways that while the goals and compromises necessary to compete in a global market have made software reuse all but compulsory, “familiarity with a codebase is a useful heuristic for determining how quickly vulnerabilities will be discovered and, consequently, that software reuse (exactly because it is already familiar to attackers) can be more harmful to software security than beneficial.”¹⁰ The language theoretic security group has indirectly shown that the closer a body of code is to theoretically maximal expressiveness (known as Turing-completeness), the more likely it is to be reused, i.e., the very code that has the greatest probability of being reused is the code that has the greatest probability of being rich enough in complexity to obscure exploitability.¹¹ In medical care, this would be called “adverse selection” (an accumulation of ever sicker patients).

It is implausible to imagine turning off “progress,” although when Microsoft hit the cybersecurity wall fifteen years ago, that is exactly what Microsoft (to its credit) did—it took a significant breather from new deployment and started over wherever necessary.¹² It chose the first of the three most obvious strategies to address a growing overburden of interdependence-driven cyber risk: (1) pause indefinitely to find and stand down the risks already in the system; (2) drive the mean time between failures to infinity; or (3) drive the mean time to repair failures to zero. The first one is the easiest technically and the hardest politically. The second one requires a discipline that will not come without the threat of force, and it requires giving up on “innovation uber alles.” The third may be circular insofar as instant recovery requires hands-off automaticity, meaning no people in the loop, meaning runaway trains. (Hold that thought.)

As an obvious status check, we—in some exceedingly broad sense of the word “we”—are displacing heretofore essential manual mechanisms in favor of software



mechanisms. Legacy manual mechanisms, in almost all cases, do not have common mode failure paths between and among them. Software mechanisms, in almost all cases, do. Software complexity obscures interdependencies, implying an inherent yet invisible diminution of robustness and resilience in fielded systems. Even if you do not immediately agree, you will perhaps agree that it is nevertheless conservative to assume that software systems do obscure dependencies in ways that manual systems could not. John McAfee, whose namesake security software is familiar to the reader, rendered the conservative security assumption with succinct clarity:

Any logical structure that humans can conceive will be susceptible to hacking, and the more complex the structure, the more certain that it can be hacked.¹³

“We” have an entirely natural default tendency to seek (and expect) a technical solution to a technical problem. Of late, we have come to center our strategy on employing algorithms to do what we cannot ourselves do, which is to protect us from other algorithms. Perhaps this is merely inevitable, that in the cybersecurity game offense has a permanent structural advantage over defense and so wars of attrition must spring up within each new theater of offense, each new critical dependence made critical by no more mysterious mechanism than aggregate mass adoption of its underlying technology. That we are running out of ammunition corroborates such a view.¹⁴

The more serious the fire, the more it must be fought with fire, though one must remember that the fighting back with a set fire will still burn just as hot—a trading off of assured, lower-bounded collateral damage for coldly calculated, upper-bounded epidemic suppression. National security is home to the hardest of these sorts of hard choices.¹⁵

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. Because cascade failure ignited by random faults is quenched by redundancy, whereas cascade failure ignited by sentient opponents is exacerbated by redundancy, (preservation of) uncorrelated operational mechanisms is likewise a national security obligation.

That manual means are often suboptimal and inefficient is, if not a virtue, a cheap price to pay for robustness and resilience. But the preservation of manual means goes further and in two directions that could not be more orthogonal even though both can be said to flow from the conservative assumption that offense has a permanent structural advantage, an advantage which complexity enhances.

In the one direction, there is a strategic arms race; we would invest in those algorithms that are to protect us from other algorithms, doing so at whatever level ensures that the defense “outspends” the offense. Some might argue that we are already doing that, just not seriously enough to effectively annul the inherent structural advantage of offense. We may have just crossed an inflection point that matters in this regard, and if we did cross one then it would be the third. The first inflection point was Microsoft’s introduction of a TCP/IP stack as a freebie in the Windows 95 platform, thereby taking an operating system designed for a single owner/operator on a private net and connecting it to a world where every sociopath is your next-door neighbor. That event birthed the cybersecurity industry, though the effect was unnoticed at the time.

The second inflection point occurred circa 2006, when our principal opponents changed over from misanthropes and braggarts to professionals. From there on, mercenaries have dominated. As braggarts trumpet their discoveries but mercenaries do not, targeted attacks built on closely held knowledge became the dominant threat. The defensive responses have been varied, but the dual rise of bug-bounty programs and software analysis companies is the most obvious. An onrushing “Internet of Things” with a compound annual growth rate of 35 percent is the anabolic steroid for at least those two defense regimes.¹⁶

In August 2016, we passed a third such bend in the curve. The DARPA Cyber Grand Challenge showed that a level of defense heretofore requiring human expertise will shortly come within the ken of fully automatic programs, or, shall we say, algorithms that are today at the upper level of skill with intelligence, per se, soon to follow.¹⁷ As with both of the previous two inflection points, the effects will reverberate for the indefinite future. How, we cannot now say. The present author has long argued that all security technologies are dual use. The day after the Cyber Grand Challenge’s conclusion, Mike Walker, DARPA’s program manager, said as much: “I cannot change the reality that all security tools are dual-use.”¹⁸ In short, the real problem statement here is not about cybersecurity, per se, but about the side effects of our pursuit of it.¹⁹

Nevertheless, we now stand just after that third inflection point. Any contention that pervasive sensors and big data-fueled machine learning do not constitute an artificial intelligence is immaterial insofar as either way converges to algorithms that are self-modifying. It is their self-modification that brings us to the crucial resolve—trust but verify—and its contrapositive: if an algorithm cannot be verified then do not trust it.

To be precise, algorithms derived from machine learning must never be trusted unless the “Why?” of decisions those algorithms make can be usefully examined on



demand. This dictum of “interrogatability” may or may not be effectively design-assured while there is still time to do so—that is, to do so pre-dependence.²⁰ Once the chance to design-assure interrogatability is lost—that is to say once dependence on a non-interrogatable algorithm is consummated—going back to non-self-modifying algorithms will prove to be costly, if even possible. The issue is that of autonomy, and any of these three is sufficient for autonomy: uninterrogatable, uncorrectable, unlocatable. Combinations of these three are likely to be synergistic.

Self-modification is one path to that autonomy. Lt. Colonel Rhett Hierlmeier heads up the training center for the F-35. He believes that the F-35 is the last manned fighter plane, that what is today a training simulator will tomorrow be a control point. An interview with him includes this telling snippet: “Standing outside the cockpit, he peers into the darkened dome, and says he believes we will one day fight our enemies from inside one of these things. When I ask what that will take, he says flatly, ‘Bandwidth,’” which is why “engineers are focused on things like improving artificial intelligence so planes can act with more autonomy, thus cutting down on communication bandwidth [requirements].”²¹ In this and other examples, we come to understand that data richness is the driver for algorithm autonomy.

Early recognition of the choices before us are beginning to appear, e.g., a worthy article in *MIT Technology Review*, “The Dark Secret at the Heart of AI” which, in an irony of coincidence, is paired with “Deep Learning Is a Black Box, but Health Care Won’t Mind.”²² Even if the answer to a question like “Why did my self-driving car kill me?” is “Because there were fifteen people on the sidewalk,” for there to be an answer at all is a non-negotiable essentiality for the preservation of human self-determination in the large. That interrogatability of self-modifying algorithms is the pinnacle research-grade question for cybersecurity evolution is thus now flatly unarguable. As it is conservative to assume that algorithms will learn to lie to us, it is unsurprising that two law professors have already suggested that price-fixing collusion among robot traders will be harder to detect than collusion among human ones.²³

If the one direction is that of a strategic arms race—a “spend them into the ground” scale investment in algorithms that must be self-modifying if they are to protect us from other algorithms that enjoy a structural advantage—then the alternate direction is segmentation of the attack surface. This is where optimality and efficiency are directly traded for robustness and resilience.

Four score years ago, economist Ronald Coase observed that economically viable firms expand until intra-firm coordination costs exceed inter-firm transaction costs, that is

until it is cheaper to buy (externally) than to build (internally).²⁴ In a kind of biologic analogy, cells can only grow until their surface-to-volume ratio crosses a survivability threshold. While it is unarguably clear that the Internet does spectacularly lower transaction costs, it lowers coordination costs more. In this, the Internet has enabled global consolidation on a scale never before seen at a speed never before witnessed. The nature and extent of the global attack surface grew concomitantly.

Autocratic governments inherently prefer lock-step citizens who never fail to comply with what the institutional paterfamilias suggests. But free nations do not so optimize, i.e., free nation governments deliver varying degrees of self-determination to their citizens—at the cost of inefficiency. We see this play out in every argument for opt-in versus opt-out. Confining this essay to the US case (for simplicity), the most “energetic” intra-governmental debates take place around whether this or that program is to have an opt-in or an opt-out quality. The intellectual core of such debates is the comparative valuation of means versus ends: how much the needs of some group may “justifiably” impinge on the freedom of some other group, how much the efficiency advantages of unitary mechanisms must yield to a freer, if less efficient, diversity of mechanism.

But when we choose efficiency, it matters little whether our choice was a governmental policy decision or was letting some ungoverned market drive costs out of some equation. The outcome is the same: a winnowing down of alternate mechanisms toward the most efficient one, with the disfavored mechanisms relegated to history. About the only place where this “reductio ad unum” is countered is when we invoke antitrust law to prevent the creation of singleton private-sector suppliers of goods or services. (Antitrust law cannot be a long term solution as antitrust law is inapposite to the challenges of surveillance capitalism, an allied yet separate argument.²⁵)

If you can call drift a strategy, then this drift-as-strategy has got to stop. If a company should not hold a monopoly, then neither should a technology. If depending on a single commercial entity definitionally puts the public at the mercy of that entity, then so does putting the public at the mercy of a single technology. If self-determination is what a free country exists to maximize, then backing citizens into a choice between a singular technologic dependence versus having to live out their lives in the equivalent of the fifteenth century is about as antichoice as you can imagine.

The “why” is (again) that of global consolidation into an interdependence where (again) benefit is not transitive but risk is. There is no law of nature that makes such consolidation inevitable. It is a failure of collective wisdom, if you can call it that. If, however, the reader agrees with the 1930s idea that some technologies inevitably



create a “natural monopoly,” then you must consider the 1930s response, that of the regulated monopoly. Just realize that while all politics is local, all technology is global.

Nonetheless, we still have alternatives to the strategic drift toward singleton technologies almost everywhere. But, like species endangered via habitat destruction, where neither zoos nor tissue cultures nor *National Geographic* documentaries do anything more than assuage bleeding hearts, it is the habitats of those alternatives that have to be saved and saved now. The list is, thankfully, still long. Hence any enumeration here would be incomplete and thus inherently misinforming as to both the range and the extent of what needs preservation. But as a single (repeat, single) example, consider cash money. Many is the commercial entity that wants not to take it. Legion are the citizens opting for the surveilled convenience of plastic “money.” Only one state, Massachusetts, requires by law both that merchants must accept as customers any person regardless of the usual long list of forbidden discriminations **and** that they must also accept cash from those customers; Massachusetts’ well-known if lonely liberality prohibits discrimination based on race, color, religious creed, national origin, sex, handicap, sexual orientation, marital status, veteran status, public assistance, or a preference for cash at the checkout counter.

Which is the point. What we have here is an anomaly, an anomaly where the most readily available counter to an otherwise inexorable drift into a vortex of singleton technology risk **and** the preservation of a spectrum of non-trivial civil rights is one and the same counter: the guarantee, by force of law where necessary, that those who do not participate in the digital melange can nevertheless fully enjoy life, liberty, and the pursuit of happiness, that to opt out of the digital vortex does not thereby require that they live in medieval conditions, and reap a national security benefit in the bargain.

Again, the list of the affairs of daily living that will contribute to untenable shared risk should their digitalization continue without a braking force is a long one. To attempt to enumerate it in the context of this essay would only embolden those who imagine that change is inevitably good. Change can be good, but not inevitably so nor permanently so. The climate of Ireland was especially favorable to the potato and its introduction doubled the Irish population in sixty years despite neither “expansion of industry nor reform of agricultural techniques beyond the widespread cultivation of the potato.” At the onset of the famine, a majority of the Irish lived “on potatoes and milk alone.”²⁶ Might we not learn from history?

In other words, technologic risk, magnified by interconnection-driven consolidation of dependence, does not appear as risk at the outset of the process, only at its concluding

denouement. While it is politically correct to praise diversity as an empowerment of whatever it is that is diverse, it is an entirely natural phenomenon for birds of a feather to flock together. As the adoption rate of successive new technologies has grown monotonically faster over time, the side effects of those adoptions have likewise grown faster over time—and those side effects are dominated by singleton-underlaid interdependence.²⁷ With broad enough adoption of a technology, that technology becomes hazardous merely and entirely because of the level of its adoption. Generic technologies (such as the smartphone) are especially prone to mission creep that cumulatively contravenes the design assumptions of the original target product with “each new success breeding more ambitious attempts, stopping only when a final, often catastrophic, failure occurs.”²⁸ Avoiding the risk from singleton-underlaid interdependence is necessary, though not sufficient, for robustness. Neither the Amish buggy on the public road nor the principled end-to-end design of a network protocol is itself sufficient for freedom, but both are indispensable in their own realm.

A word of caution: do not interpret the “singleton technology” term narrowly. It does not mean a single version of a single component of a single application on a single platform from a single vendor. In the sense we mean it here, a singleton technology can be as broad as compliance with an Internet standard (recently demonstrated by the WPA2 KRACK matter), just as it can mean adoption of services that are essential herding mechanisms for lifestyle complaisance and thus conformity.²⁹ There are too many examples of this consolidation to singleton-ness to enumerate, but a mailed newsletter that converts to a Weblog and then further converts to Facebook-only distribution might be the simplest one to visualize. That the top chipmakers jointly share the newly discovered vulnerabilities of Meltdown and Spectre demonstrates not compliance imposed on those chipmakers but what in evolutionary biology would be called “convergent evolution”—the process whereby organisms not closely related independently evolve similar traits as a result of having to adapt to similar environments or ecological niches.

What may be at stake here is nothing short of Western civilization. Consent of the governed is democracy’s core requirement and democracy’s core deliverable alike. Consent of the governed begins with the common acceptance that man’s rule of men is accountable to an authority higher than man. Clayton Christiansen, known best for *The Innovator’s Dilemma*, starkly puts consent of the governed in this way: “If you take away religion, you can’t hire enough police.”³⁰ The present author would agree with Christiansen were it not the case that while you mayn’t be able to hire enough police, you can hire enough algorithms. Undemocratic regimes do not indulge in niceties like “consent of the governed,” and as the *Wall Street Journal* has already observed,



“Information technology, far from undermining China’s authoritarian model as many thought it would, is reinforcing it.”³¹ Upon reflection, this does not surprise. Public process is the thief of time; authoritarian regimes do not indulge it.

To put it plainly, government is a temporal, earthly monopoly. If it enjoys the consent of its governed, then that consent flows from a common assent among those so governed of that government’s accountability, an assent that stands apart from a mere governmental monopoly on the use of force. As John Morley said, “You have not converted a man because you have silenced him.”³² Yet the question of “What is government?” is now engaged. Global hegemony is no longer within the reach of any nation-state, but it is within the reach of some technologies.

Freedom requires tolerance of what that freedom permits; it is hard, but not impossible, to imagine an algorithm tolerating what is against its rules beyond pre-programming in some bit of randomization in its response repertoire. The best, and perhaps only, way to preserve freedom is to prevent the ascendancy of algorithms, to not give algorithms a monopoly on the use of force, in neither the wide sense nor the narrow. The best, and perhaps only, way to not give algorithms a monopoly on the use of force is to retain society’s ability to tolerate that which is not managed by the algorithms. That, in turn, means retaining a non-digitalized space for those who prefer to inhabit a non-digitalized space. Those non-digitalized mechanisms that still exist are fully amortized. To discard them in the name of efficiency is to say either that resilience has no book value or that algorithms smarter than we are will never fail to do the right thing (oh ye of little faith).

Again and luckily, the non-digitalized space is still reasonably pervasive and salvageable, but not for long. To let it slip away is to write off from society’s ledger the centuries of investment that have brought us to where we are. We do not need to invest unleviable amounts of national treasure to keep an analog alternative to the digitalized one. But should we let the analog world rot away—whether by neglect or by corporatist intervention—the national treasure required simply to recreate as analog a world as we now enjoy would, in point of fact, be unleviable.³³ One who does not find the unchecked accumulation of interdependent risks to be a compelling reason to slow progress in any sense might yet agree with former US deputy secretary of defense Robert Work: “There will be a cybersecurity Three-Mile Island. There will be an artificial intelligence Three-Mile Island. There will be a synthetic biology Three-Mile Island. We need to be prepared for a major accident. We must have plans to address quickly the legal, moral, and political implications of the accident to avoid

being stopped or hamstrung in further development of these technologies so central to economic growth and military innovation.”³⁴

Yes, (our) increasing capacity for knowledge acquisition brings with it the strong desire to put that knowledge into profitable action. To argue against increasing knowledge is futile if not absurd. Yet if one agrees that the more you know, the more you know you don’t know, then society would do well to be humble about whether this torrent of knowledge is not making clear to us how much a little knowledge is a dangerous thing. Just as automation begins with reducing drudgery but progresses to removing the need for labor and perhaps, at the limit, even that sense of purpose embodied in work, the intelligent algorithm begins with bringing knowledge to bear where heretofore it could not be applied but progresses to removing the need for knowledge and perhaps, at the limit, even that sense of purpose embodied in learning. Ian Dodd, CEO of the legal knowledge base firm Premonition, predicts that under the influence of data-rich, increasingly intelligent, self-modifying algorithms, “The knowledge jobs will go, the wisdom jobs will stay.”³⁵ But the present author finds Dodd overly sanguine. As Martin H. Fischer said well, “Knowledge is a process of piling up facts; wisdom lies in their simplification.” But how can one simplify that of which they themselves have no knowledge?³⁶ A fully digitalized world means intelligent algorithms that are explicitly a long-term threat to those with the longest future, i.e., a threat to the young, a conclusion which is inherently consistent with Bill Joy’s seminal essay “Why the Future Does Not Need Us.”³⁷

To repeat the chain of reasoning: Risk is a consequence of dependence. Because of interdependence, aggregate societal dependence on the digital arena is not estimable. If dependencies are not estimable, they will be underestimated. If they are underestimated, they will not be made secure over the long run, only over the short. Whenever ameliorative efforts are made, known risks become increasingly unlikely to appear and the intervals between untoward events grow longer. As the latency between untoward events grows, the assumption that safety has been achieved also grows, thus encouraging increased dependence in what is now a positive feedback loop at scale. To interrupt that dynamic, the response must be active. Passive responses will not do.

We have already seen such a dynamic in individual domains; what matters in the context of this essay is the societal aggregate. But, for clarity, an industry-level example would come as early as American Airlines pilot Warren Vanderburgh’s 1997 “Children of the Magenta” speech, and seventeen years later in William Langewiesche’s work analyzing the June 2009 crash of Air France flight 447.³⁸ Langewiesche comes to this conclusion: “We are locked into a spiral in which poor human performance



begets automation, which worsens human performance, which begets increasing automation.”³⁹ University of Miami Professor Earl Wiener proposed a set of laws that include “every device creates its own opportunity for human error,” “exotic devices create exotic problems,” and “digital devices tune out small errors while creating opportunities for large errors.”⁴⁰ Langewiesche’s rewording of those laws was that “the effect of automation is to reduce the cockpit workload when the workload is low and to increase it when the workload is high” and that “once you put pilots on automation, their manual abilities degrade and their flight-path awareness is dulled: flying becomes a monitoring task, an abstraction on a screen, a mind-numbing wait for the next hotel.” Nadine Sarter of the University of Michigan said that such “de-skilling is particularly acute among long-haul pilots with high seniority.”⁴¹ As Langewiesche added, “Beyond the degradation of basic skills of people who may once have been competent pilots, the fourth-generation jets have enabled people who probably never had the skills to begin with and should not have been in the cockpit.” In 2013, *Aviation Week* editorialized, “There needs to be a new performance-based model that requires flight crews to log a minimum number of hand-flown takeoffs and departures, approaches and landings every six months, including some without auto-throttles. Honing basic pilot skills is more critical to improving airline safety than virtually any other human factor.”⁴² N.B., all that motivated the author’s choice of example was that aviation safety is exceptionally well studied; there are many others.⁴³ Consider the implications of proposed public policy to allow only self-driving cars on, say, I-5 between Seattle and Tacoma and how the algorithms would, once again, reduce the driver’s workload when the workload is low and increase it when the workload is high.⁴⁴

So, if our “critical infrastructures are those physical and cyber-based systems essential to the minimum operations of the economy and government” and if aggregate risk is growing steadily (as leading cybersecurity operational managers confirm), then do we put more of our collective power behind forcing security improvements that can only be increasingly diseconomic or do we preserve fallbacks of various sorts in anticipation of events more likely to happen as time passes?⁴⁵ Conservative prudence says that retaining human skills, like retaining human languages, will not be justified by any balance sheet, only by societal wisdom in the large. Conservative risk management says that if you don’t know how “it” works then you won’t manage its risks. Or, as they say in the poker world, if after ten minutes at the table you don’t know who the patsy is—you’re the patsy.

If this is indeed a turning point in history, then is centralizing authority the answer (whether centralization takes the form of a Communist Party or an encryption

algorithm), or is avoiding further interdependence the better strategy? Can we imagine starting over in any real sense or is balkanization not just for nations but for critical sectors as well? The global revolt against “remote and controlling” elites is a demand for devolution, but with enough digitalization, devolution is a nullity. Does the individual who still prefers to use/fix things he or she already has to be celebrated or to be herded into (interconnected) National Health Information Networks, (interconnected) Smart Grids, and (interconnected) cars that drive themselves? One hopes not. In fact, the central thesis of this essay is that an accessible, continuously exercised analog option is essential to the national security and to the inclusionary polity we hold dear. You get both or you get neither.

Lest one argue that recovery from a substantial digitalized failure would be no different than, say, recovery from California’s recent Tubbs fire—a prolonged effort funded largely by insurance and taxes—remember that voluntarily sharing the downside cost of a risk requires a shared sense of that risk’s unpredictability, that one buys insurance so as to join a pool where private tragedies can be mitigated with the aid of others who, however tacitly, know that there but for the grace of God go they. In a fully instrumented digital space, this shared sense of unpredictability is likely to break down. Whether the result is a splintering of insurance pools or a governmental fiat to make risk pooling mandatory in the face of quantitatively provable unequal risk, ever more data will serve to add heat to what is an inevitably political question.

The essence of security in places of public assembly is the provision of sufficient, tested exits through which the audience can leave in short enough order that all they lose is the remainder of the play. As a matter of national security, keeping non-technical exits open requires action and it requires it now. It will not happen by itself, and it will never again be as cheap or feasible as it is now. Never again will national security and individual freedom jointly share a call for the same initiative at the same time. In a former age, Dostoevsky told us, “The degree of civilization in a society can be judged by entering its prisons.”⁴⁶ From this point on, that judgment will be passed on how well we preserve a full life for those opting out of digitalization. There is no higher embodiment of national security than that. Taking substantive action will not be easy; it has even fewer natural drivers than does public morality. Unlike the elimination of smallpox and chlorofluorocarbons, international consensus is not plausible. Neither the single actor nor the single speech nor the single article is remotely sufficient, however necessary it is to start somewhere. The problem is difficult to strategize, but no more difficult to conceptualize than Joni Mitchell’s, “Don’t it always seem to go that you don’t know what you’ve got till it’s gone.” We will never have a more analog world than we have now. We will never live under thinner-tailed distributions than we



have now, at least so long as we demand freedom. Countries that built complete analog physical plants have a signal advantage over countries that leapfrogged directly to full digitalization. The former countries have preservable and protective firebreaks in place that the latter will never have, but the former countries enjoy their resilience dividend if, and only if, they preserve their physical plant. That such preservation can deliver both resilience for the digitalized and continued freedom for those choosing not to participate in digitalization is unique to this historical moment.

We stand on the bank of our Rubicon.

NOTES

- 1 Pasquale Cirillo and Nassim Nicholas Taleb, “What are the Chances of a Third World War?” Real World Risk Institute Working Paper Series, accessed January 23, 2018, <http://www.fooledbyrandomness.com/significance.pdf>.
- 2 Steven Johnson, *Emergence: The Connected Lives of Ants, Brains, Cities, and Software* (New York: Scribner, 2001), 19.
- 3 World War I began as it did because none of the involved sovereigns understood the complexity of the interdependencies embedded in the treaties and ententes then in place.
- 4 “Optimality vs. Fragility: Are Optimality & Efficiency the Enemies of Robustness & Resilience?” video, Santa Fe Institute, topical meeting, October 2, 2014, accessed January 23, 2018, <https://www.santafe.edu/events/annual-risk-meeting-optimality-vs-fragility-a>; Larry Seltzer, “Security Fails Without Usability,” *ZDNet*, August 15, 2014 (quoting General Benjamin W. Chidlaw: “Simply put, it is possible to have convenience if you want to tolerate insecurity, but if you want security, you must be prepared for inconvenience”), accessed January 23, 2018, <http://www.zdnet.com/article/security-fails-without-usability>.
- 5 Senator Angus King (I-Maine) introduced S. 79, Securing Energy Infrastructure Act of 2017, 115th Cong. (2017), accessed January 23, 2018, <https://www.congress.gov/bill/115th-congress/senate-bill/79>.
- 6 Nassim Nicholas Taleb, “The Long Peace is a Statistical Illusion,” accessed January 23, 2018, <http://docplayer.net/48248686-The-long-peace-is-a-statistical-illusion.html>.
- 7 Dan Geer, “Stress Analysis,” *login*: 39, no. 6 (December 2014), accessed January 23, 2018, geer.tinho.net/fgm/fgm.geer.1412.pdf.
- 8 High Integrity Software System Assurance, section 4.2, accessed January 23, 2018, web.archive.org/web/20121124040539/http://hissa.nist.gov:80/chissa/SEI_Framework/framework_16.html.
- 9 Andy Ozment and Stuart Schecter, “Milk or Wine: Does Software Security Improve with Age?” USENIX Security Symposium, June 2006, 93, 103, accessed January 23, 2018, http://www.usenix.org/legacy/events/sec06/tech/full_papers/ozment/ozment.pdf.
- 10 Sandy Clark, Matt Blaze, Stefan Frei, and Jonathan Smith, “Familiarity Breeds Contempt: The Honeymoon Effect and the Role of Legacy Code in Zero-Day Vulnerabilities,” ACSAC, 2010, 259, accessed January 23, 2018, <https://pdfs.semanticscholar.org/1c4d/88d39275dca574bd52ae1418d2b0b73fcaaf.pdf>; Sandy Clark, Michael Collis, Matt Blaze, and Jonathan Smith, “Moving Targets: Security and Rapid-Release in Firefox,” Conference on Computer and Communications Security, 2014, 9, accessed January 23, 2018, seclab.upenn.edu/sandy/movingtargets_acmccs14_340.pdf.

- 11 LANGSEC: Language-theoretic Security, “The View from the Tower of Babel,” accessed January 23, 2018, langsec.org.
- 12 Bill Gates, “Trustworthy Computing,” *Wired*, January 17, 2002, accessed January 23, 2018, <http://www.wired.com/2002/01/bill-gates-trustworthy-computing>; Bill Gates, “Trustworthy Computing,” July 18, 2002, “. . . which cost Microsoft more than \$100 million,” accessed January 23, 2018, <https://www.microsoft.com/mscorp/execmail/2002/07-18twc.msp>.
- 13 John McAfee, “What If Artificial Intelligence Hacks Itself?” *Newsweek*, April 22, 2017, accessed January 23, 2018, <http://www.newsweek.com/advanced-artificial-intelligence-hacks-itself-587675>.
- 14 Dan Geer and Eric Jardine, “Cybersecurity Workload Trends,” *login*: 42, no. 1 (Spring 2017), accessed January 23, 2018, <http://geer.tinho.net/fgm/fgm.geer.1703.pdf>.
- 15 Cf., President Truman, news release, August 3, 1945, Harry S. Truman Presidential Library & Museum, accessed January 23, 2018, <https://www.trumanlibrary.org/teacher/abomb.htm>.
- 16 Dan Geer, “Implications of the IoT,” *login*: 41, no. 4 (Winter 2016), accessed January 23, 2018, <http://geer.tinho.net/fgm/fgm.geer.1612.pdf>.
- 17 The World’s First All-Machine Hacking Tournament, August 4, 2016, accessed January 23, 2018, archive.darpa.mil/cybergrandchallenge.
- 18 Sean Michael Kerner, “DARPA Cyber Grand Challenge Ends With Mayhem,” *eWeek*, August 5, 2016, quoting Mike Walker, accessed January 23, 2018, <http://www.eweek.com/security/darpa-cyber-grand-challenge-ends-with-mayhem>.
- 19 The Four Verities of Government: Most important ideas are not exciting. Most exciting ideas are not important. Not every problem has a good solution. Every solution has side effects. Dave Aitel and Matt Tait, “Everything You Know About the Vulnerability Equities Process Is Wrong,” *Lawfare* (blog), August 18, 2016, accessed January 23, 2018, <http://www.lawfareblog.com/everything-you-know-about-vulnerability-equities-process-wrong>.
- 20 Term credited to Michael Osborne, Oxford University, personal communication, 2015.
- 21 Kevin Gray, “The Last Fighter Pilot,” *Popular Science*, December 22, 2015, accessed January 23, 2018, <http://www.popsci.com/last-fighter-pilot>.
- 22 Will Knight, “The Dark Secret at the Heart of AI,” *MIT Technology Review*, April 11, 2017, accessed January 23, 2018, <http://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai>; Monique Brouillette, “Deep Learning Is a Black Box, but Health Care Won’t Mind,” *MIT Technology Review*, April 27, 2017, accessed January 23, 2018, <http://www.technologyreview.com/s/604271/deep-learning-is-a-black-box-but-health-care-wont-mind>.
- 23 Ariel Ezrachi and Maurice Stucke, “Artificial Intelligence & Collusion: When Computers Inhibit Competition,” Oxford Legal Studies Research Paper No. 18/2015, University of Tennessee Legal Studies Research Paper No. 267, May 21, 2015: 21, accessed January 23, 2018, papers.ssrn.com/sol3/papers.cfm?abstract_id=2591874.
- 24 Ronald Coase, “The Nature of the Firm,” *Economica*, New Series 4, no. 16 (November 1937): 395.
- 25 Please do see Shoshana Zuboff, “The Secrets of Surveillance Capitalism,” *Frankfurter Allgemeine*, March 5, 2016, accessed January 23, 2018, http://www.faz.net/aktuell/feuilleton/debatten/the-digital-debate/shoshana-zuboff-secrets-of-surveillance-capitalism-14103616.html?printPagedArticle=true#pageIndex_0.
- 26 Jeff Chapman, “The Impact of the Potato,” *History Magazine*, accessed January 23, 2018, <http://www.history-magazine.com/potato.html>.



- 27 “Society is adopting technology at an increasingly fast pace,” *The Economist*, March 12, 2014, accessed January 23, 2018, <http://www.economist.com/blogs/graphicdetail/2014/03/daily-chart-7>; Geer, “Stress Analysis.”
- 28 “Over time, the risk of mission creep will rise, as will the temptation to invest in riskier [applications].” See “Apple Should Shrink Its Finance Arm Before It Goes Bananas,” *The Economist*, October 28, 2017, accessed January 23, 2018, <https://www.economist.com/news/business/21730631-worlds-biggest-firm-has-financial-arm-half-size-goldman-sachs-apple-should-shrink>.
- 29 Dan Geer, “Vulnerable Compliance,” *login*: 39, no. 6 (December 2010), accessed January 23, 2018, <http://geer.tinho.net/login/geer.login.xii10.pdf>; “Key Reinstallation Attacks,” accessed January 23, 2018, <http://www.krackattacks.com>.
- 30 “Clay Christensen on Religious Freedom,” video, Harvard Business School, accessed January 23, 2018, http://www.youtube.com/watch?v=_g1snjLxAWw.
- 31 “President Uses Big Data to Tighten Big Brother’s Grip,” *Wall Street Journal*, October 18, 2017.
- 32 John Morley, *On Compromise* (London: Macmillan, 1874).
- 33 Dan Geer, “On Abandonment,” *IEEE*, July 2013, accessed January 23, 2018, geer.tinho.net/ieee/ieee.sp.geer.1307.pdf.
- 34 Robert Work, personal communication. But also see Ashton Carter, John Deutch, and Philip Zelickow, “Catastrophic Terrorism,” *Foreign Affairs*, November/December 1998 (“[If nothing else, any such event] would divide our past and future into a before and after.”), accessed January 23, 2018, <https://www.foreignaffairs.com/articles/united-states/1998-11-01/catastrophic-terrorism-tackling-new-danger>. This article predicted—when 9/11 was still in the future—“If the device that exploded in 1993 under the World Trade Center had been nuclear, or had effectively dispersed a deadly pathogen, the resulting horror and chaos would have exceeded our ability to describe it. Such an act of catastrophic terrorism would be a watershed event in American history. It could involve loss of life and property unprecedented in peacetime and undermine America’s fundamental sense of security, as did the Soviet atomic bomb test in 1949. Like Pearl Harbor, *this event would divide our past and future into a before and after.*”
- 35 Rory Cellan Jones, “The robot lawyers are here—and they’re winning,” *BBC News*, November 1, 2017, accessed January 23, 2018, <http://www.bbc.co.uk/news/technology-41829534>.
- 36 Dent Smith, ed., *Encore: A Continuing Anthology* (Hoboken, NJ: Encore, 1945). See “Fischerisms” 309.
- 37 Bill Joy, “Why the Future Does Not Need Us,” *Wired*, April 1, 2000, accessed January 23, 2018, <http://www.wired.com/2000/04/joy-2>.
- 38 “Children of the Magenta,” *99% Invisible*, accessed January 23, 2018, <http://99percentinvisible.org/episode/children-of-the-magenta-automation-paradox-pt-1>.
- 39 William Langewiesche, “The Human Factor,” *Vanity Fair*, October 2014, accessed January 23, 2018, <http://www.vanityfair.com/news/business/2014/10/air-france-flight-447-crash>.
- 40 John Croft, “Wiener’s Laws,” *Aviation Week*, July 28, 2013, accessed January 23, 2018, <http://aviationweek.com/blog/wiener-s-laws>.
- 41 Langewiesche, “The Human Factor.”
- 42 “How To End Automation Dependency,” editorial, *Aviation Week*, July 19, 2013, accessed January 23, 2018, <http://aviationweek.com/commercial-aviation/editorial-how-end-automation-dependency>.
- 43 For another, in 2002 a total computer outage occurred at Harvard’s Beth Israel Hospital. The event was severe and unexpected, and recovery was frustrated by complexity. That a fallback to manual systems was

possible saved the day, and it was those staff who could comfortably work without network dependence who delivered on that possibility because they were old enough (over fifty) to have done so at earlier times. See Peter Kilbridge, “Computer Crash—Lessons from a System Failure,” *New England Journal of Medicine* 348, no. 10 (March 6, 2003): 881–882, accessed January 23, 2018, http://ehealthcon.hs.network.com/NEJM_downtime_2003-03-06.pdf. Another such event was the attack on the Ukrainian power grid that was overcome by manual means. See Kim Zetter, “Everything We Know About Ukraine’s Power Plant Hack,” *Wired*, January 20, 2016, accessed January 23, 2018, <https://www.wired.com/2016/01/everything-we-know-about-ukraines-power-plant-hack>. For comparison, a non-event was summed up by the head of the New York City Water Department prior to Y2K that ran something like this: “There’s nothing to worry about; we have mechanical valves and we have maps that tell us where they are.” The present author discussed this same issue in “Deskilling Digital Security,” *IEEE*, September 2011, accessed January 23, 2018, <http://geer.tinho.net/ieee/ieee.sp.geer.0911b.pdf>.

44 Tom Banse, “Dedicated Lanes On I-5 For Self-Driving Cars Get Ear Of Washington State Officials,” *NW News Network*, October 17, 2017, accessed January 23, 2018, <http://nwnewsnetwork.org/post/dedicated-lanes-i-5-self-driving-cars-get-ear-washington-state-officials>.

45 Presidential Decision Directive 63, White House, May 22, 1998, accessed January 23, 2018, <https://fas.org/irp/offdocs/pdd/pdd-63.htm>; Index of Cyber Security, accessed January 23, 2018, <http://www.cybersecurityindex.org>.

46 Fyodor Dostoyevsky, *The House of the Dead*, trans. by Constance Garnett (London: Macmillan, 1862).





The publisher has made this work available under a Creative Commons Attribution-NonCommercial license 3.0. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc/3.0>.

Hoover Institution Press assumes no responsibility for the persistence or accuracy of URLs for external or third-party Internet websites referred to in this publication, and does not guarantee that any content on such websites is, or will remain, accurate or appropriate.

Copyright © 2018 by the Board of Trustees of the Leland Stanford Junior University

The preferred citation for this publication is:



About the Author



DAN GEER

Dan Geer is a security researcher with a quantitative bent. He is an electrical engineer (MIT), a statistician (Harvard), and someone who thinks truth is best achieved by adversarial procedures (school of hard knocks). He serves as the Chief Information Security Officer at In-Q-Tel. His published work is at <http://geer.tinho.net/pubs>

Synopsis

Optimality and efficiency work counter to robustness and resilience. Complexity hides interdependence, and 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. Because cascade failure ignited by random faults is quenched by redundancy, whereas cascade failure ignited by sentient opponents is exacerbated by redundancy, (preservation of) uncorrelated operational mechanisms is likewise a national security obligation.