A HOOVER INSTITUTION ESSAY

Flat Light

DATA PROTECTION FOR THE DISORIENTED, FROM POLICY TO PRACTICE

ANDREW BURT & DANIEL E. GEER, JR.

Aegis Series Paper No. 1816

Flat light is the state of disorientation, feared among pilots, in which all visual references are lost. The effects of flat light "completely obscure features of the terrain, creating an inability to distinguish distances and closure rates. As a result of this reflected light, [flat light] can give pilots the illusion that they are ascending or descending when they may actually be flying level."[1]

This is the state of information security today.

Attack surfaces have expanded beyond any organization's ability to understand, much less defend against, potential adverse events. Common interdependencies, once assumed secure, are not, rendering entire protocols, infrastructures, and even hardware devices susceptible to exploitation.

So large is the deluge of potential security threats that a new phrase has entered the lexicon for information security professionals: "alert fatigue."[2] One 2015 study, focused on malware triaging efforts at over 600 US organizations, found an average of 17,000 alerts generated per week, with only 4 percent of such alerts ever investigated.[3] And that's just malware alerts. The information we have at our disposal about our vulnerabilities does little in the way of mitigating them.

The problem, then, for information security practitioners and policymakers—including government officials, lawyers, and privacy personnel—is one of bearing. When you don't know where you're going, all directions are equally useless. We simply do not know what to focus on, how to spend our energy, what precise regulation is called for, or how to significantly disincentivize would-be attackers.

But this state of affairs has not always been the case.

While under siege since its earliest days, the world of information security has always had reference points—or ground truths—that, like physical features in a landscape, have served as guides to practitioners and policymakers alike. These reference points, which we detail below, provided at least a modicum of bearing to those engaged in data protection.

As the aggregate state of information security has deteriorated over time, however, features of this landscape have eroded under the pressure of a changing environment, rendering past reference points either unhelpful (at best) or disinformative (at worst).

Flat light is now upon us.

We aim, in this paper, both to explain how we arrived at this situation, at least in part, and to suggest a path forward.

I. What Was True Then

Reference points create structure among objects moving through space and time. With the proper reference points, a pilot can understand where she is heading and at what speed. Such references might include features in a landscape like ridges or canyons, large bodies of water, major buildings, radio beacons, and more.

The world of information security has had its own reference points, whether explicit or implicit, that helped practitioners and policymakers navigate their choices. We aim here to describe what this world looked like when these reference points were clear or, at the very least, not so clouded.[4]

For those who wish to skip the history lesson, you may proceed to Section II, in which we outline the reference points that currently define—or are in the process of defining—the cybersecurity landscape today.

> ***What is "cybersecurity"?***
>
> Any discussion of this subject, with as broad and an admittedly ambitious scope as ours, must contend with this simple fact: nobody knows what "cybersecurity" means in practice. The concept has, until recently, been used as the sum total of the reference points we overview in Section I. The problem, as we note throughout this paper, is that many such concepts have become obsolete.
>
> As such, "cybersecurity" is a collective noun in transition, destined to become the new sum total of the still-evolving reference points we describe in Section II, yet not having reached its final target. While we use the term throughout the paper, we are fully aware of its shortcomings, and ask the reader to be so aware as well.

Cyberspace as "space"

We begin with the notion that ***activities committed in "cyberspace" were informatively analogous to real world counterparts.***[5] Perhaps one of the earliest reference points in the world of information security, this idea is literally embedded in the United States Code. The Computer Fraud and Abuse Act (CFAA) of 1986, for example, imposes criminal penalties on cybercrime premised upon the perceived parallels between physical acts and digital

activities.[6] The central thesis of the CFAA is that the boundary between legal and illegal acts in cyberspace rests upon unauthorized *access,* just as torts like trespass are premised upon crossing a physical threshold between property that is "mine" by items that are "yours."[7]

The very notion of cyberspace as an extension of physical space helped to extend traditional notions of privacy into the digital domain. Indeed, ***the long-running distinction between being observed and being identified*** has provided the very basis for digital notions of privacy. These notions of privacy trace their way back to the publication of the 1890 article "The Right to Privacy" in the *Harvard Law Review* by future Supreme Court Justice Louis Brandeis and attorney Samuel Warren.[8] Being *observed* in a public setting, the theory went, constituted one type of activity, having little import to an individual's "privacy." But it was a qualitatively different type of activity to be *identified* in public, at least without some type of prior consent.[9]

Granting cyberspace the veneer of physicality also implied that ***security and privacy were two fundamentally different activities.*** Privacy, per Warren and Brandeis, consisted of the right "to be let alone." That right manifested as an implicitly legal activity within organizations, largely left to lawyers or compliance personnel, who focused on how users expected to be observed as they conducted their digital activities or generated new data. As a result, privacy laws have been largely written for and carried out by compliance personnel, rather than more technical information security practitioners.

Information security practice, on the other hand, was borne from the existence of attackers: sentient opponents who sought to penetrate networks via vulnerabilities waiting to be exploited. In the world of privacy, the phrase "hands on keyboard" meant writing a memo; in the world of security, it meant activities within a command screen.

Thinking of cyberspace as an analog to physical space has felt both comforting and natural in that it minimizes the profound, sometimes vexing, differences between acts taking place within computer networks and actions in the corporeal world. But as we detail in Section II, this framing happens to no longer be true.

Failure is not an option.

With cyberspace as an analog to physical space, it followed that data could be protected in ways that paralleled our own notions of physical safety. In the physical world, security frequently means the complete absence of harm, at least aspirationally. Physical security measures don't simply aim to minimize harm to physical objects (like the human body); they aim to generate protection that is as all-encompassing as possible. This approach was extended to the world of cyberspace, where the goal has been to ***eliminate failure, not merely to minimize it.***[10]

As such, practitioners and policy makers alike both sought ways to drive the mean time between failures (MTBF) as high as possible. The higher the MTBF, the less frequently

failures affect your network. ***True success was, at least implicitly, defined as an MTBF of infinity.*** This idea relied upon the assumption that a failure could and would be recognized as being a failure.

Through that lens emerged the so-called CIA triad, composed of the concepts of confidentiality, integrity, and availability. This triad—while far from the whole story—has guided data protection efforts since their earliest days, mirroring, in effect, how one might go about protecting information in a book or a private journal.[11] Within the CIA triad, ***confidentiality has always been prioritized*** as the first among equals.[12]

Like a lock on the outside of a personal diary, the dictates of confidentiality were aimed at ensuring data were only viewed by individuals who had a key, or by those who were granted access by someone who did. The idea that the information inside the journal might be corrupted (addressed by the notion of integrity), or that it might not be accessible when needed (encompassed by availability), has generally been secondary to the aim of keeping that information private.

When issues of integrity did arise, another reference point held that ***binary ideas of "correctness" dictate how data should be protected.*** Data were either correct or they weren't; they had been tampered with or their integrity remained fully intact. Efforts to keep data protected were either successful or they had failed. Keeping track of this binary status is what determined whether or how well data had been protected. Indeed, after confidentiality, ensuring that data were untampered with has ranked as the next most important goal of both information security professionals and information security policies.

The avoid-failure-at-all-costs ideal led to the reference point that ***diversity is a sufficient protection from attackers.*** The foundational idea came from nature: monocultures maximize risk, where a single infectious agent can trigger an extinction. Avoiding homogeneity—by encouraging diversity—was a sufficient protection from the worst outcomes of the worst attacks. Given enough diversity, *something* would survive.

In all such cases, failures in information security were viewed as binary—and to be avoided. This possibility of total avoidance was a comforting notion that, while once possible, is no longer feasible, even with automation.

Ability, and stability, to predict

Past reference points about information security were also based upon temporal assumptions relating to our ability to predict what types of systems, or what specific locations within such systems, required increased levels of protection. There was—or at least there appeared to be—a definable baseline stability underlying the systems we sought to protect. Ordinary behavior was reflexively good. Anomalies were reflexively bad.

As a result, a huge number of information security practices became premised on the assumption that ***it is feasible to determine what is normal á priori.*** This is most particularly seen in its extension to systems like "critical infrastructure," which US administrations have sought to coherently define for decades, as well as to systems within organizational networks.[13] Indeed, the very idea that systems can be designated as "critical" and their "normal" pattern of operation cataloged forms the backbone of information security efforts in countries throughout the world.[14]

Along the same lines, it was understood that we could ***define the nexus of a network, or its most important physical endpoints, by assessing its physical facilities.*** The physicality of these facilities, which we could plot on a map, helped determine what types of data needed to be protected, and where.

But technologic trends—the massive adoption of networked technologies, at ever smaller units of deployment, feeding ever expanding data to ever cheaper compute power, on networks defined solely by their software—have eroded all underlying bases of stability. The rate of change created by our adoption of networked technologies has undermined our very ability to understand what it is we are adopting. The end result is combinatoric explosion—the number of possible modes of interaction is increasing exponentially as ever more components are installed.[15]

We may yet be able to explain our networks today, but we cannot make long-term predictions about them. "Normal" versus "abnormal" may have been a useful concept. But at high enough rates of change, the distinction becomes a nullity.

(The Illusion of) Control

Because computers run programs, and programs are fundamentally straightforward constructions, it has long been a reference point that ***better programming (and programmers) would necessarily make for better outcomes,*** be it defense against hostile actors or resilience to known stimuli. The list of idealized solutions connected to this line of thought was long: if only programmers would pay attention to security from the get-go; if only management were to concretely prioritize security for their programmers; if only customers cared whether the app they so desperately wanted were written by programmers paying attention to security as prioritized by their management; if only . . . ad infinitum.[16]

But complexity hides interdependencies that have rendered these solutions less meaningful. Every major security breakdown of the last decade, for example, has had more than one cause.[17] With the rise—and increasing utilization—of machine learning, that complexity is set to increase.[18] There will, in short, always exist more than one point where better programmers, building for better customers, as prioritized by better management, will still not be able to anticipate the path to an exploitable vulnerability.

Indeed, a dual reference point we have long relied upon—from consumers to programmers to policymakers—is the hope that ***security can be made composable,*** meaning that with skillful technique the interconnect of two secure programs would be inherently secure. That assumption is false, and proven so regularly.[19]

The illusion of control now extends to the very concept of ownership as well. Ownership—the basic idea of applying property rights in cyberspace—extends less and less to both computing platforms and our data. It is not new, for example, that computing platforms are essentially "rented" through license agreements, rather than purchased outright by consumers. But computing environments are now composed of a far-reaching series of owners. Rarely can the end user—from an individual to an enterprise—claim much, if any, stake in that environment.[20]

To express this point more directly, we once assumed that ***consumers could dictate the terms of their service.*** In the past, one might have been able to theoretically assert about a software product, "I bought it. It's mine. I tell it what to do." No longer.[21]

II. A Shift in Truths

The above set of reference points, whether acknowledged or even recognized, allowed both policymakers and information security practitioners to prioritize their activities and to orient their energies in relation to the threats they believed they faced. Like the visibility of physical landmarks to navigators, each reference point gave us the ability to assert order amid a gale of challenges.

We do not contend that this order never existed, simply that it no longer does.[22]

In its stead, new reference points have arisen to take the place of what we documented above. We highlight ten such reference points below.

1. Actions in cyberspace decreasingly resemble activities in the physical world.

It is said that the practice of law is the search for analogies. The notion of a cyber "space" analogized two common ideas: that the boundaries between networks physically resembled, and could be thought of as, the boundaries between physical bodies; and that the data we use and seek to protect can usefully be thought of as something akin to physical objects.

The more we move to cloud-based infrastructure, however, the more these conceptions become nonsensical. Our data live less and less on our own local devices and are instead stored remotely in massive data centers, which depend on parallelized computing and virtual machines.[23]

Thirty years ago, we took our data to where the computing was—a university's central computing facility, perhaps. With the innovation of desktop computing, it became possible to move the computing to where the data were, and individuals gained easy access to that processing. Then came virtualization—which creates multiple simulated environments within a single hardware system—and the distal end became once again a display tool, while the data have gone back to where the computing is (only now we can't tell where, exactly, "where" is). The next oscillation has already begun: computing is going back to where the data reside—in the sensor fabric.

As illustrated in the case *In re Search Warrant No. 16-960-M-01 to Google,* thinking about data as residing in one place, or even as one coherent collection of objects, distorts reality. In that case, Google refused to produce records for the US government because the company breaks down its communications into multiple components, which are dynamically stored in different geographical locations. As a result, the company argued that it lacked the capability to determine the precise location of any communication from a user at a specific point in time.[24]

While we leave the substance and logic of Google's arguments to the courts, we use the case to illustrate the growing mismatch between legal conceptions of data and technical ones.[25]

2. Twentieth-century notions of privacy no longer apply to our activities.

What were once clearly physical activities have now become digitized—that is, turned into and stored as data—in nearly every realm of our existence. Our locations are constantly mapped and stored in relation to cellular towers; our communications are now conducted almost exclusively in digital form; our cars have essentially become software with engines and wheels. The list goes on. The very word "surveillance"—attentive, continuous, purposive observation— has become nearly meaningless in its ubiquity.

All the data we generate—while seemingly trivial in isolation—yield potentially intimate insights about who we are and what we do when those data are aggregated (or "fused"). Thinking about our data physically gives the false impression that the combination of our data is additive, such that each new correlation adds "1" to some pre-existing sum.

It is not.

The combinatory power of our data is not additive, it is multiplicative, like a compound interest calculation, making it nearly impossible for individuals to understand how their data will be used or what insight they may, in the aggregate, deliver over time. For the same reason data scientists can't predict what their conclusions will be before training a model on a data set, we cannot expect individuals to understand what facts and insights they "give up" while generating data. For this reason, privacy frameworks predicated on user

consent—in which we rely on users to make meaningful, informed decisions about their own privacy—are destined to fail in the aggregate, precisely because no user can know the true value of her data at the point it is generated.[26]

There is, for example, the famous (publicized) case of the retailer Target knowing that a teenager was pregnant before her family did, based simply on her shopping habits, which Target simply "observed" in public, as it were.[27] In 2016, a group of researchers claimed to be able to identify, with nearly 90 percent accuracy, whether an individual had committed a crime simply based on patterns in facial features.[28] Such examples abound, and more are coming.[29]

The ability to draw deeply invasive conclusions about us from disparate, seemingly inconsequential data grows by the day. Meanwhile, it has become impossible to refrain from generating these types of data, at least if one is to nominally participate in modern society.

What Warren and Brandeis once called the right "to be let alone" is, quite simply, no longer feasible. There can be no "alone" when so much of our daily activities generate unintentionally—and uncontrollably—sensitive data. New conceptions of privacy must be formulated if we are to salvage what's left of the information we seek to protect. We suggest a basis for such conceptions in Section III.

3. Privacy and security are converging.

As traditional notions of privacy have eroded, so too has the distinction between privacy and security. New laws focused on data now blend privacy protections with security mandates, like the European Union's General Data Protection Regulation (GDPR) or China's Cybersecurity Law. The GDPR, for example, mandates a host of privacy provisions, ranging from what it calls privacy "by design and by default" to specific measures focused on the more traditional CIA triad.[30] That legislation is among the most recent examples of the blending of privacy and security into one field with a more singular, intuitively simpler end-goal: the control of data.

At its root, the problem statement for both digital rights management, which seeks to enforce copyright and use protection on digital media, and for privacy enhancing technologies, which we discuss below, is one and the same: How do you control the use of information you own at a distance across both space and time?[31]

Phrased another way: What, exactly, are we protecting?

In the worlds of both data security and privacy, the answer is increasingly the *misuse* of our data, either by current adversaries or by those we trust at present but may not in the future. And because we are not, and cannot be, fully aware of the potential value of all the data we

generate, we never know precisely what it is we entrust to others when we hand them our data. Anyone, from a privacy perspective, can become an adversary, given enough time.

Indeed, the world in which organizations possessed data they didn't understand and used those data for reasons into which they had limited insight is fading. Instead, new data-focused legislation, combined with increasing concern over data collection, is moving us toward a world where data protection is focused on preserving the rights and expectations of the subjects who generated the data *and* on protecting those data from potential adversaries.

That world cannot come fast enough.

From a technical standpoint, given the increasing adoption of machine learning techniques, this convergence makes sense. Machine learning models make decisions based on recognizing patterns in large volumes of data. These models are susceptible to new types of malicious activity that raise both privacy and security concerns, frequently at the same time. By gaining insight into how such models make their classifications, for example, researchers have been able to conduct attacks that simultaneously undermine both the privacy and security of those models.[32]

What used to be a clear distinction between privacy and security, quite simply, no longer exists. What we call "privacy" and "security" are now best and jointly described as "data protection."[33]

4. Failures cannot be avoided.

Failures—be they lapses in confidentiality, integrity, or availability—must now be assumed, both for policymakers and for information security professionals alike. As the interdependence of systems whose designers are unaware of each other grows, unpredictable events necessarily occur. Like asking, "Will California suffer a bad earthquake?" the answer to whether unforeseen, adverse events will take place in environments of increasing complexity will always be affirmative—probabilistic events occur eventually. The question, then, becomes *when* (or *how soon*) such events will occur.[34]

Approaches to risk management in areas outside of information security are instructive. Even in the high-stakes world of medicine or aviation where mistakes can lead to loss of life, failure tolerance is not set to zero.

We should then ask: "How many failures is the right number of failures?"

The answer cannot be "zero," nor should it be. No failures means over-spending, both from a resources perspective and in opportunity costs. Failures are a primary component in how

we learn. Failures should, when approached systematically and when their risks are managed carefully, be central to the cost of evolution. No failure, no advance.[35]

It is from this reference point that our next arises.

5. The goal of information security consists of driving the mean time to repair, or MTTR, to zero.

If failures are inevitable, the aim of information security cannot—and should not—be to drive the MTBF to infinity. Rather, if failure is a given, what matters is the length of the time between *failure* and *repair.*[36]

Take the case of Chaos Monkey, a service created by Netflix to randomly terminate its infrastructure systems.[37] Given the company's reliance on infrastructure susceptible to random failures, Netflix engineers created a service to ensure failures occurred, and therefore built resiliency into their operating model by default. The cause of such failures was, of course, internal, rather than the exploitation of vulnerabilities by a third party, but the end result was the same: a focus on repair rather than prevention.

We are witnessing a shift, then, from MTBF, which was premised on avoiding failures, to MTTR (mean time to repair), which emphasizes our ability to react to and remedy adverse events when they occur.

Central to the utility of repair is the ability to detect. And this means that failure detection—and the strict avoidance of silent failures—will be among the most significant tasks in data protection, if it is not already. We noted the existence of "alert fatigue" at the outset of this paper. We fear the future might hold the reverse in store, where the onset of failure is slow and delayed, arising from causes unknown. We may well face a world in which the material cause of any failure cannot be meaningfully extracted from networks too complex for our understanding.

Experience with large software systems tells us how supremely dangerous failure opacity can be over time. There are multiple reasons why, but the simplest to describe derives directly from the versatility of the system and the expressiveness of its many interfaces. If that expressiveness permits subtlety, then rule-sets expressed in the system tend to grow over time because operators of that system will tend to avoid analyzing whatever subtleties have accumulated. They will instead blindly add special-case allowances and constraints, so as to avoid breaking whatever rules came before.

In very little time, the possibility of comprehensively addressing problems diminishes and we are left with attempts at symptomatic relief.[38] This problem will only be magnified

as machine learning is adopted, where the source of the models—the data on which the models were trained—may not only be difficult to understand or to access, they may be wholly unavailable.

6. Integrity has eclipsed confidentiality in the CIA triad.

Given all these changes, we must also reevaluate the way we approach the CIA triad. Specifically, if we assume that failures in confidentiality, integrity, or availability will inevitably occur, confidentiality cannot have the priority it once held. What now matters most is integrity—namely, understanding and proving which data can be trusted after failure has occurred.

We return to the increasing prominence of machine learning as one reason for this shift. In a world where machine learning is deployed more pervasively, the importance of integrity increases. Machine learning models are shaped by the data they train on. In simpler terms, they eat data for breakfast, lunch, and supper too, all the while obeying the rule that "you are what you eat."[39] Feed the models cleverly corrupted data, and the models are vulnerable in new, unpredictable ways.

Understanding whether—and, if so, how—data were corrupted becomes an increasingly critical task.

7. Tracking data provenance, and not simply correctness, is now paramount.

From the growing importance of data integrity arises the significance of data provenance, which we loosely define as the ability to trace and record the origins and movement of data across databases. Provenance is fast becoming central to data protection.[40]

Indeed, we submit that provenance is, in fact, now a core part of integrity itself. If we do not know where the data came from or what processing they underwent post-collection, we cannot protect their integrity. The same goes for the algorithms those data have been used to create. In many ways, the realities of machine learning mean that we do not have source code, we have source data. And as sources proliferate, we no longer have tests for correctness, we have only probabilities.[41]

8. Diversity without adaptability does not provide sufficient protection against attackers.

Diversity has been deemed a central component in protecting systems from single-cause failures that, in the physical world, lead to devastating outcomes (species extinction, famine, pandemics, etc.).[42] The reason is that only diversity is inherently able to hold off common-mode failure.[43]

In the world of data protection, however, a range of developments is already beginning to decrease the value of diversity. Automated exploitation tools, for example, render the value of diversity less than it once was. The growth of networks—and the related interdependencies between complex systems—is making diversity harder to sustain and to manage.[44]

In the physical world, the attacker must commit a perfect crime and the police have all the time in the world to unravel the mystery. In the digital world, the police must frequently craft the perfect defense while the attacker has all the time at her disposal to find a single flaw in that defense. This asymmetry, borne from the speed and distance from which attacks take place, underpins the demand to have consistent, airtight security. In practice, the phrase "consistent and airtight" just about means "all alike," and all alike is the perfect setting for attacks to come without warning. If you are lucky, the black swan doesn't come on your watch.

In other words, in addition to diversity, adaptability among systems is now critical. This is especially true amid the universe of systems making up the so-called Internet of Things. Such adaptability must allow for continuous updates and monitoring, enabling these systems, which constitute an ever-greater proportion of our attack surface, to become nonstatic and, therefore, less susceptible to predictive methods of attack. The only alternative is that those devices and systems have a built-in, known-in-advance, finite lifetime. Either way, the combination of unfixable and immortal is untenable.

This points us to a fork in the road, which all system designers now face. On the one side is provable, defect-free correctness in code, followed by brutally rigorous change control; on the other side is so-called moving target defense, rapid release, and DevOps, all premised upon agility in development and nimble reaction to unforeseen events as they arise. The former is increasingly achievable for well-defined stand-alone systems but remains out of reach for volatile interdependencies of the sorts found in the types of software systems supported only by advertising.[45] The latter only works when the rate of change is ratcheted upward so that the opponent cannot reverse-engineer the most recent change before another is thrust upon her.[46] These two approaches are antithetical. Yet both are supported by sound scientific research and best practices.

All of which is to say that diversity and adaptability are both achievable—just not all at once.

9. Critical infrastructure is based on adoption rates.

What makes something critical infrastructure? Adoption. Adoption is the gateway drug to criticality. When enough people depend on something in cyberspace, that something is made critical.[47]

As has been demonstrated repeatedly, the rate of adoption for new technologies has grown over the last century.[48] This means that essential aspects of society come into their position of criticality with ever less lead time. The law—and its practitioners—cannot keep up.

A sustained and accelerating flux of change puts defenders at a long-term structural disadvantage compared to offenders. Long-term predictions upon which planning depends become less trustworthy. Self-modifying algorithms, to pick just one technology, make this dynamic abundantly clear.[49] In environments containing such models, we may simply never know what our "attack surface" is. The very concept of attack surface becomes an outmoded term the minute we no longer know how an algorithm is adapting to new data. Rate of change is one basis of offense's structural dominance over defense. As such, public policy faces the conundrum of whether to slow the rate of change or to cede control of its side effects.

This reality stands in stark contrast to the approaches we described above, in which policy makers could actively assert what infrastructure within a network was "critical." The changes occurring within the data protection landscape simply cannot be fully mapped out in advance. Neither can they be predicted. We have become disoriented precisely because we seek stability from an environment that can no longer produce it.

It is from these insights that we draw our final reference point.

10. Your nexus is no longer based on the physical location of your facilities but is based instead upon the location of your users.

If we cannot predict what constitutes critical infrastructure in advance, does the very idea of critical nodes in a network remain relevant?

The answer is yes. But the twist is that what makes these nodes critical is the adoption rate of the network's users themselves, as we noted above. No longer does a network's physicality translate to meaningful insight. It is who is using the network, and the location of the users, that matters, not the location of the resources those users use.

In practice, this means an Internet defined by the locations of its users: in other words, a balkanized Internet.[50] The Chinese government has long understood this truth, basing its entire cybersecurity strategy on destination control.[51] The European Union's data localization requirements, as exhibited in the GDPR, are a close second.[52]

We predict that the Internet will be "geocoded" with a precision similar to geocoding the mobile phone fleet. Paired with personalization based on surveillance, the Internet will appear differently depending on the combined location and identity of each user.[53] Of course, if cellular phones become the dominant access point to the Internet (as looks all

but inevitable), then a geocoded Internet is simply a matter of database management—a problem that we've long known how to manage, and to manage well.[54]

III. A Path Forward

What was once a complex, though navigable, landscape has now become increasingly onerous. So, from a practical standpoint, what to do?

We suggest a few answers below.[55]

Restrict data based on use

To begin with, ***purpose-based restrictions on data must be a central component to protecting privacy,*** or what's left of it**.** Restrictions on use must become a core component of data protection laws around the world. This isn't to say restrictions on collection—which seek to limit what types of data organizations can collect and when—are completely useless. They are simply *more* useless every day.

Because we cannot understand the insights that massive amounts of data will yield á priori, we cannot fully protect data as they are generated. Indeed, in a world of ubiquitous connected devices, we can't even understand which data we're generating or how we're generating them. And because we cannot stop creating this ever-growing mass of data, attempting to protect our privacy rights at the point of collection is futile. As Bruce Schneier puts it, "We are getting better, but we are getting worse faster."[56] We might make progress, but any gains will quickly be undone.

In practice, this means that data handling restrictions must be enforced throughout the entire data lifecycle, forcing organizations in a position to collect data to justify exactly *why* they are utilizing specific datasets and legally bounding the potential reasons for that use.[57] To be meaningful, rather than simply cosmetic, such restrictions must be enforced and monitored by a blend of information security professionals, lawyers, and data engineers.

Laws must use intent to define crimes in cyberspace.

Laws directed at digital activities must focus on the intent underlying the action, in addition to the means of action. The CFAA in the United States, the country's major cybercrime law, is perhaps the biggest culprit in this area, creating adverse ramifications for activities in cyberspace far beyond the United States' borders. The CFAA defines cybercrimes as consisting of "unauthorized access" into a computer or network, which, as we noted above, is a direct result of the misconception that cyberspace is analogous to the physical world.

But what if that unauthorized access occurred in the process of researching vulnerabilities so that they can be fixed? Or what if a criminal organization sells software designed to enable unauthorized access but doesn't conduct that activity itself?

Under the CFAA, the answers to both of these questions appear to be the wrong answers. Security research is criminalized, pure and simple, if that research results in access that a network operator has not explicitly authorized, even if the intent of the research is *to benefit the security of that network and its users*. As a result, the law has had a serious and deeply chilling effect in the information security community.

And because the law overemphasizes the importance of access, activity that should obviously be illegal—like the rental of botnets or other programs that promote unauthorized access—is, at the very least, arguably in the clear. If I license you a program that allows you to hack into another computer or network, it is not clear that I have done anything wrong, because I haven't done the actual "hacking." This is like enshrining in law that the government can only go after low-ranking Mafiosi, when it is the high-ranking ones who do the most damage. The CFAA's focus on acts alone, rather than intent, retards the pursuit of criminal justice in cyberspace, as some officials within the US Department of Justice have themselves noted.[58]

We expect the current flaws of the CFAA to worsen as machine learning algorithms become ever more prominent. The studies we referenced above, in which cleverly manufactured input data were fed into machine learning models to generate incorrect predictions, do not clearly constitute unauthorized access and are not clearly illegal under the CFAA. And yet these types of adversarial activities could directly result in human injury if, say, deployed against image recognition systems in self-driving cars to manipulate their classifications, to cite just one example.

The fix to these problems is clear: laws governing crimes in cyberspace should be reshaped to focus on the intent underlying the malicious activity. If an academic researcher can prove or certify the intent of the activity is for research, and if she or he did not cause any specific or demonstrable harm, that activity should not be illegal. Conversely, if actions undertaken in cyberspace are intended to create or promote malicious activity, those actions should be illegal. The onus should be on the actor to prove intent, so extreme caution would still reign supreme.

This approach is not without challenges, but there is sufficient legal precedent governing intent, or mens rea, that we do not believe this legal construct would be prohibitive from a procedural standpoint. Coupling mens rea with actus reus, the act committed, is already the central component of much of criminal law. It should be extended to crimes governing computing activities as well.[59]

Embrace digital identification

The drawbacks of identities in cyberspace are enormous. There are potential threats to freedom of speech, the loss of the ability to conduct activities anonymously, and many more. And yet, despite these downsides, there is simply no way around the formalization

of identification online, which will allow organizations—and, yes, governments—to prove that an online persona corresponds to an offline individual during certain activities.[60]

Many such attempts are in process, such as the controversial Aadhaar system in India and the Estonian system of online IDs.[61] While these systems and others have significant shortcomings, the underlying goal of these efforts should not be dismissed out of hand.

If we seek to assert control over our data—which is, we contend, the ultimate goal of privacy efforts—then ***we must be able to formally track which data belong to whom.*** In other words, if you want to exercise your right to be forgotten, you first have to prove who you are.

Positives outcomes will surely arise from the formalization of IDs online, from faster, more efficient disbursement of government and other services to better monitoring of fraud and other malicious behavior. These benefits should be welcomed by all.

But that's not why we make the case for online IDs here.

We make the case because, in truth, an online identification system *already* exists. To argue otherwise is to fight the last war. Take as a given the ease with which our information can be de-anonymized. Take as a given that social media is a system of informants. Take as a given that what is observable by technical means defeats any attempt to be let alone in the sense of Brandeis and Warren.

This identification system all but exists in practice. It merely hasn't been formalized, meaning that we already suffer its downsides without reaping its benefits. The wrong actors can identify us in cyberspace, should they choose to; the right actors, whom we might seek to act on our behalf, don't yet have that option. This fact is neither pleasant nor fashionable, but it is nonetheless true.

We offer no recommendations into precisely *how* to offer a system of digital IDs. Volumes have been written on the subject elsewhere, by opponents and proponents of the idea alike. Our main point is merely that this system, to the extent it doesn't already exist, is inevitable. Those who oppose digital identification efforts can best exert their influence by attempting to shape that system and prevent its abuses, before those who don't share these concerns begin to.[62]

Retain the analog

Individuals and organizations must have the effective capacity to forgo digital services altogether, meaning that ***analog alternatives to digital activities must be preserved,*** even if we cannot preserve past conceptions of privacy altogether. It cannot be that individuals are forced to choose between the extremes of either "all your data belong to someone else" or "go live in the fifteenth century." Indeed, even the latter choice, while far from

ideal, may no longer exist absent substantive moves on the part of governments, which we encourage.

As is currently being seriously discussed in Sweden, for example, the fraction of the citizenry choosing to be cashless is beginning to depress cash volume to the point that those who do not wish to use "digital money" will be left with no money at all.[63] As the world saw in Puerto Rico after Hurricane Maria and Florida after Hurricane Michael, cashless out-of-choice and cashless out-of-disaster are all but indistinguishable.[64]

But that is only one example. The entire issue of analog alternatives to digital services is at once a matter of avoiding cascade failure—something to which analog systems are largely immune, but which digital systems are all but universally vulnerable—and a matter of human rights, if you believe that digital rights are human rights, as we do.[65] The need for analog alternatives is then both a moral issue and a national security obligation if we are to have resiliency in our networks and our infrastructure.[66]

In practice, the preservation of analog alternatives requires that those means be used, not left to gather dust in the hope than when they are needed they will still work. This requires a base load—a body of use and users that keeps the analog working.

And while all politics may be local, all technology is global. We make this point to stress that solutions to these problems, taking place in various jurisdictions throughout the world, are all ultimately interrelated. The Internet may in fact be balkanizing, but neither local solutions nor local issues can be fully separated from the global challenges we confront.

Mandate sustainability

Because failure is now a given, and because diversity is no longer sufficient protection from would-be adversaries, the sustainability of our networks and infrastructure must be mandatory. By "sustainable," we mean resilient to changing conditions over time. We cannot allow connected devices to become static and unadaptable indefinitely, as is occurring with the devices that make up the Internet of Things. ***If a device has an IP address, it must either accept updates or have a finite lifetime.*** This is an either-or design choice that must be required by law.[67]

Sustainability may be an over-used word, but it applies to the minimal need for failure tolerance, which must be based on the ability to receive updates and patches or a designated end of life for Internet-connected devices. The growing impact of the "CAP theorem," which describes the trade-offs among consistency, availability, and partition tolerance, cannot be ignored.[68]

Indeed, the range and scale of the problems we face are already evident in the world of mobile devices where, as the US Federal Trade Commission noted in a 2018 report, these

devices are "not receiving the security patches they need to protect them from critical vulnerabilities."[69] This report wasn't based on information from small-scale or little-known manufacturers that we might expect to constitute the Internet of Things. It was based upon information collected from eight of the largest mobile device makers on the planet: Apple, Blackberry, Google, HTC, LG, Microsoft, Motorola, and Samsung. If these device makers can't collectively manage security updates, how can we expect better from the makers of "smart" lightbulbs, toasters, fitness trackers, and more?

The answer: we cannot.

The lesson: sustainability must be legally mandated and not left to a market that prioritizes time to deployment. Alternatives to sustainability cannot, and should not, be available under the law.

More data is better

By now it's become commonplace to suggest that more data equal better solutions, as is the default mentality in so much of the world of engineering. On the other hand, no data equal no solutions—or no good solutions—which is where we oftentimes find ourselves in the world of data protection.

Indeed, due to the increasing unpredictability of the data protection landscape, we now require real-time data for our decisions and for our policies. For this reason, data relating to information security—the data about our data—need to be more easily accessible to organizations in both the private and public spheres.[70]

A few legislative attempts have already been made with these goals in mind. One such example is the Cybersecurity Information Sharing Act of 2015, or CISA, which created a mechanism for private actors to share and receive threat information with and from the United States government through the Department of Homeland Security.[71] That attempt, while laudable, has failed.[72]

Under CISA, there is no reward for sharing such information, and too many risks in doing so, to incentivize enough organizations to submit and receive the volume of data required to create the workable, collective situational awareness to properly understand our environment. In CISA's wake, a few private-sector information-sharing groups have sprouted to accomplish the same ends, though none are yet on a large enough scale to create the type of holistic—or easily and freely accessible—information landscape we need.

The European Union's GDPR also requires data sharing from businesses to consumers in what it calls a right to data "portability" which would, in theory, allow users to easily move their data between the organizations that collect it, just as a cellphone number can currently be transferred between carriers in the United States. We don't believe

this attempt will fall quite as flat as CISA's efforts. But the right to access, retain, and transfer data between organizations is destined to run into countless technical and procedural hurdles, precisely because the complexity of IT environments is increasing. Again, this legislative attempt is laudable, but likely to fall far short of creating the type of environment where data are both easily accessible and inherently protected by organizations and individuals alike.

Instead, what's needed is a focus on new technologies and techniques that make data sharing faster *and* privacy enhancing, such as the group of approaches referred to as privacy enhancing technologies (PETs), briefly referenced above and frequently promoted by privacy advocates.[73] Such technologies include differential privacy, homomorphic encryption, federated learning, and more.

In a broad sense, PETs use mathematical techniques to modulate the trade-off between utility, security, and privacy—a trade-off that all too frequently minimizes security and privacy when left to the pressures of the real world, where speed to market trumps all else. PETs can allow data to be pooled and shared across organizations without significantly increasing the risks in using those data.[74]

While we refrain from directly recommending any one particular technology here, our central point is that ***novel approaches to data sharing must be embraced*** if we are to overcome the problem that we currently face, where none but a few global technology giants have the resources to fully utilize all the data we generate. Those data must form the basis of our decisions, in as close to real time as possible, if we are to have any hope of acting intelligently with the choices we face.

IV. Conclusion

Concrete conclusions require order. Our precise thesis is that the existing order is changing—disintegrating, even—in the wake of the new reference points we set forth in Section II. From these emerging truths we draw a handful of recommendations, ranging from changing existing laws on security and privacy, to creating new ways to identify and share individuals' data, to ensuring the survival of analog alternatives to digital activities.

Despite the admittedly capacious topics we've discussed, however, our overall aim is more simplistic and can be summed up with one word: understanding.

We have lost the basic ability to understand our data protection environment. And we cannot protect what we cannot understand. As it stands, we confront a future in which failures may occur without detection, attack surfaces cannot be cataloged or quantified, and we bear more blame than our adversaries for the outcomes that follow.

This is a world dominated by flat light.

"So what should a pilot do when all visual references are lost?" asks a guide to flat light written by the Federal Aviation Administration.[75]

The answer: "Execute a 180-degree turnaround and start looking for outside references."

It is not too late to follow that advice.

NOTES

1 "Aeronautical Information Manual Study Guide for the Private Pilot" (Elite Aviation Solutions, 2013), 324, accessed November 8, 2018, https://goo.gl/74XbzC.

2 See, for example, Ryan Francis, "False Positives Still Cause Threat Alert Fatigue," *CSO Online,* May 3, 2017, accessed November 8, 2018, https://www.csoonline.com/article/3191379/data-protection/false-positives-still-cause-alert-fatigue.html.

3 "The Cost of Malware Containment," Ponemon Institute, January 2015, p. 1, accessed November 8, 2018, https://www.ponemon.org/local/upload/file/Damballa%20Malware%20Containment%20FINAL%203.pdf. Reports such as Ponemon's have been confirmed time and again. See, for example, "McAfee Labs Report Finds 93 Percent of Security Operations Center Managers Overwhelmed by Alerts and Unable to Triage Potential Threats," Intel Newsroom, December 2016, p. 6, accessed November 8, 2018, https://newsroom.intel.com/news-releases/mcafee-labs-report-finds-93-percent-of-security-operations-center-managers-overwhelmed-by-alerts-and-unable-to-triage-potential-threats.

4 Note that while we attempt to describe past reference points of import, we don't claim, or even attempt, to list all such markers or "truths" exhaustively. That is to say, there are key markers that we have left out in this paper.

5 Throughout this essay, we denote reference points and key recommendations by stating them in ***boldface italics.***

6 See 18 U.S. Code § 1030 (defining criminal activity as having "knowingly accessed a computer without authorization or exceeding authorized access").

7 We take the CFAA's focus on access—among a host of other potential digital activities—to be one among many signs of its physical framing. Indeed, the relatively short law, codified at 18 U.S. Code § 1030, makes reference to "access" no fewer than eighteen times. A less analogously physical framing, in our view, might have focused on the misuse or destruction of data.

8 Samuel D. Warren and Louis D. Brandeis, "The Right to Privacy," *Harvard Law Review* 4, no. 5 (December 15, 1890), accessed November 8, 2018, http://www.cs.cornell.edu/~shmat/courses/cs5436/warren-brandeis.pdf.

9 See Andrew Burt and Dan Geer, "The End of Privacy," *New York Times,* October 5, 2017, accessed November 8, 2018, https://www.nytimes.com/2017/10/05/opinion/privacy-rights-security-breaches.html; and Dan Geer, "The Right to Be Unobserved," *IEEE Security & Privacy,* July/August 2015: 88, accessed November 8, 2018, http://geer.tinho.net/ieee/ieee.sp.geer.1507.pdf.

10 In this case, failure was—and for the purposes of this paper, is—broadly defined as an unanticipated event with negative consequences. We readily admit that failures have been inevitable in the world of information security. Here, we simply assert that the underlying aim of security efforts has been predicated on the fundamental assumption that failures *could* be detected and therefore *should* be avoided where possible. We advocate an alternative approach to confronting failures.

11 See, for example, Eugene Troy, Ron Ross, David Ferraiolo, Eugene Bacic, and Jonathan Wood, "Perspectives and Progress on International Criteria" (paper presented at 15th National Computer Security Conference,

National Institute of Standards and Technology, Baltimore, MD, October 13–16, 1992), accessed November 8, 2018, https://csrc.nist.gov/CSRC/media/Publications/conference-paper/1992/10/13/proceedings-15th-national-computer-security-conference-1992/documents/1992-15th-NCSC-proceedings-vol-2.pdf: "It is universally recognized that government and commercial institutions rely heavily on information processing systems to meet their operational, financial, and information requirements. The integrity, availability, and confidentiality of key software systems, databases, and data networks is a major concern in all industrialized nations."

12 As we note, this description is accurate *in general*—but far from the whole story. Different actors have tailored the triad to their own priorities. In the world of financial institutions, integrity has played a larger role. In the retail space, vendors have prioritized availability. In general, however, it is confidentiality that has played the most prominent role across sectors.

13 For an illustration of this reference point in US federal policies, see the White House, "Presidential Decision Directive on Critical Infrastructure Protection," May 22, 1998, accessed November 8, 2018, https://fas.org/irp/offdocs/pdd/pdd-63.htm; Department of Homeland Security, "Homeland Security Presidential Directive 7: Critical Infrastructure Identification, Prioritization, and Protection," December 17, 2003, accessed November 8, 2018, https://www.dhs.gov/homeland-security-presidential-directive-7; and the White House, "Presidential Policy Directive on Critical Infrastructure Security and Resilience," news release, February 12, 2013, accessed November 8, 2018, https://obamawhitehouse.archives.gov/the-press-office/2013/02/12/presidential-policy-directive-critical-infrastructure-security-and-resil.

14 See, for example, China's 2017 Cybersecurity Law, which designates "critical information infrastructure," as detailed by Sara Xia in "China's New Cybersecurity Law: The 101," *China Law Blog,* June 24, 2017, accessed November 8, 2018, https://www.chinalawblog.com/2017/06/chinas-new-cybersecurity-law-the-101.html; or the European Programme for Critical Infrastructure Protection, which the European Commission overviews at https://ec.europa.eu/home-affairs/what-we-do/policies/crisis-and-terrorism/critical-infrastructure_en, accessed November 8, 2018, among many other examples.

15 We use "exponentially" here in its mathematical sense, and not in the colloquial usage as a superlative for "a lot."

16 "When you want something really bad, you will put up with a lot of flaws," to quote Stewart Butterfield, founder of Slack. To the point, this very much includes security flaws, as "security" is generally a feature added only after good consumer uptake. Stewart Butterfield, "We Don't Sell Saddles Here," February 17, 2014, accessed November 8, 2018, https://medium.com/@stewart/we-dont-sell-saddles-here-4c59524d650d.

17 There is not room here to discuss this point fully as, in fact, each major breach involved a number of essential steps. For a kind of index of such breaches, however, see Taylor Armerding, "The 17 Biggest Data Breaches of the 21st Century," *CSO Online,* January 26, 2018, accessed November 8, 2018, https://www.csoonline.com/article/2130877/data-breach/the-biggest-data-breaches-of-the-21st-century.html.

18 D. Sculley et al. provide what's perhaps the best example we've come across of complexity in an environment utilizing machine learning in "Hidden Technical Debt in Machine Learning Systems" (lecture at 29th Annual Conference on Neural Information Processing Systems, Montreal, December 7–12, 2015), accessed November 8, 2018, https://papers.nips.cc/paper/5656-hidden-technical-debt-in-machine-learning-systems.pdf.

19 As aptly stated by Bruce Schneier, "Sometimes it'll be two secure systems that, when [they] interact in a particular way, cause an insecurity." Bruce Schneier, "What 'Efail' Tells Us about Email Vulnerabilities and Disclosure," *Lawfare* (blog), May 24, 2018, accessed November 8, 2018, https://lawfareblog.com/what-efail-tells-us-about-email-vulnerabilities-and-disclosure. More simply put, we cannot predict how complex systems will interact with each other once connected or what vulnerabilities might arise, even if each system is secure in isolation.

20 Nazli Choucri and David Clark provide a great overview of this (growing) complexity, spelling out the dozens of steps and actors involved in a simple act such as the loading of a webpage. See Nazli Choucri and David D. Clark, "Integrating Cyberspace and International Relations: The Co-Evolution Dilemma," MIT Political Science

Department Research Paper No. 2012-29, November 7, 2012, p. 8–9, accessed November 8, 2018, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2178586.

21 Or such, at least, has been the fiction that courts and vendors have long operated under. Courts have, for example, long evaluated contracts between consumers and software providers by asserting that consumers have some modicum of ability to assent (or decline) the terms of service and therefore have some implied measure of control over their software environment. The larger the purchasing organization, the (presumably) larger the degree of control. While this sense of control may have been, in some cases, a fiction all along, it is no longer even remotely justifiable for the reasons we outline above. As the providers of any critical service consolidate market control, such contracts become simply contracts of adhesion.

22 "Reality is the leading cause of stress amongst those in touch with it."—Jane Wagner, "The Search for Signs of Intelligent Life in the Universe," 1977.

23 Even at the level of the individual user, Google's Chromebook exemplifies inherently off-board data storage and accounts for approximately 60 percent of the US K-12 education market. See Mary Jo Foley, "Windows PCs Gain Share In K-12 In the US, But Chromebooks Still Dominate," *ZDNet,* March 6, 2018, accessed November 8, 2018, https://www.zdnet.com/article/windows-pcs-gain-share-in-k-12-in-the-u-s-but-chromebooks-still-dominate.

24 In re Search Warrant No. 16-960-M-01 to Google, 232 F.Supp.3d 708, 712 (E.D. Pa. 2017), accessed November 8, 2018, https://www.washingtonpost.com/news/volokh-conspiracy/wp-content/uploads/sites/14/2017/02/Opinion.pdf?tid=a_inl: "Google contends that it does not currently have the capability, for all of its services, to determine the location of the data and produce that data to a human user at any particular point in time." Google's arguments were ultimately rejected.

25 As legal scholar Andrew Keane Woods notes, "There is substantial case law suggesting that courts think of data as a physical object." Andrew Keane Woods, "Against Data Exceptionalism," *Stanford Law Review* 68 (April 2016): 729, accessed November 8, 2018, http://www.stanfordlawreview.org/wp-content/uploads/sites/3/2016/04/68_Stan_L._Rev._729_-_Woods.pdf.

26 The fair information principles, or FIPs, for example, originated in the 1970s and form the basis for much of modern data privacy regulation. The FIPs are in many senses built upon the importance of notice and consent. As the privacy scholar Woodrow Hartzog has pointed out, however, "The FIPs are inadequate for the modern world." Woodrow Hartzog, "The Inadequate, Invaluable Fair Information Practices," *Maryland Law Review* 76, no. 4 (2017): 952, 956, accessed November 8, 2018, http://digitalcommons.law.umaryland.edu/cgi/viewcontent.cgi?article=3759&context=mlr. See also Helen Nissenbaum's take on this subject in Scott Berinato, "Stop Thinking About Consent: It Isn't Possible and It Isn't Right," *Harvard Business Review,* September 24, 2018, accessed November 8, 2018, https://hbr.org/2018/09/stop-thinking-about-consent-it-isnt-possible-and-it-isnt-right.

27 Charles Duhigg, "How Companies Learn Your Secrets," *New York Times,* February 16, 2012, accessed November 8, 2018, https://www.nytimes.com/2012/02/19/magazine/shopping-habits.html?pagewanted=1&_r=1&hp.

28 *MIT Technology Review* provided an overview of the study in *Emerging Technology from the arXiv,* "Neural Network Learns to Identify Criminals by Their Faces," November 22, 2016, accessed November 8, 2018, https://www.technologyreview.com/s/602955/neural-network-learns-to-identify-criminals-by-their-faces. While many cast doubt upon the actual conclusions of the study, we view its mere existence as an example of the simpler, larger trend of sensitive information being available in ways that defy immediate human intuition.

29 We find new methods on de-anonymization of authorship particularly interesting with regard to both code and natural language. Studies suggest the very act of writing—almost anything—is now a potential means towards identification. See Arvind Narayanan et al., "On the Feasibility of Internet-Scale Author Identification" (presentation at 33rd Annual IEEE Symposium on Security and Privacy, San Francisco, CA, May 20–23, 2012), accessed November 8, 2018, https://people.eecs.berkeley.edu/~dawnsong/papers/2012%20On%20the%20Feasibility%20of%20Internet-Scale%20Author%20Identification.pdf; and Aylin Caliskan et al., "When Coding

Style Survives Compilation: De-anonymizing Programmers from Executable Binaries" (paper presented at Network and Distributed Systems Security [NDSS] Symposium, February 18–21, 2018 San Diego, CA), accessed November 8, 2018, https://arxiv.org/abs/1512.08546. With regards to geospatial data and privacy, Yves-Alexandre de Montjoye, César A. Hidalgo, Michel Verleysen, and Vincent D. Blondel provide a good illustration of its limits in "Unique in the Crowd: The Privacy Bounds of Human Mobility," *Scientific Reports* 3 (2013), accessed November 8, 2018, https://www.nature.com/articles/srep01376.

30 See General Data Protection Regulation Article 25 ("Data Protection by Design and by Default") and Article 32 ("Security of Processing"), among many other examples.

31 Dan Geer, "Getting the Problem Statement Right," Harvard University, Cambridge, MA, January 22, 2004, accessed November 8, 2018, http://geer.tinho.net/geer.harvard.22i04.txt. For those who don't like the term "own," you may substitute "over which you have the controlling interest."

32 Researchers, for example, have fooled image detection classifiers into labelling turtles as rifles or baseballs as espresso, presenting potentially severe security challenges. The consequences of this type of attack on an autonomous vehicle, or on nearly any robotic device, could literally prove fatal. So, too, have researchers been able to learn enough about these models to understand the original data they were trained on, including personally identifying information about individuals contained in the raw data. See "Fooling Neural Networks in the Physical World with 3D Adversarial Objects," LabSix, October 31, 2017, accessed November 8, 2018, https://www.labsix.org/physical-objects-that-fool-neural-nets; and Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov, "Membership Inference Attacks Against Machine Learning Models" (paper presented at 38th Annual IEEE Symposium on Security and Privacy, San Jose, CA, May 22–24, 2017), accessed November 8, 2018, https://www.cs.cornell.edu/~shmat/shmat_oak17.pdf.

33 A trend we readily embrace, as evidenced by the subtitle of this paper.

34 One sure sign that failure is to be absolutely prevented is choosing to have no plan for mitigating the occurrence of that failure. Any claim of "that can't happen because we won't let it" had better be backed up with a kind of zeal that is vanishingly rare and incredibly difficult.

35 Worse by far is suppressing the knowledge of failure. See, for example, Dan Geer and Richard Danzig, "Mutual Dependence Demands Mutual Sharing," *IEEE Security & Privacy* 87 (January/February 2017), accessed November 8, 2018, http://geer.tinho.net/ieee/ieee.sp.geer.1701.pdf.

36 It goes without saying that where redundancy can obviate the need for repair, redundancy can (and should) be prioritized.

37 See Chaos Monkey, GitHub, accessed November 8, 2018, https://github.com/Netflix/SimianArmy/wiki/Chaos-Monkey. For further reading on the general subject of failures contributing to resiliency, see Nassim Nicholas Taleb's *Antifragile: Things that Gain from Disorder* (New York: Random House, 2012).

38 See the longer quote from Don Davis in Dan Geer, "Application Security Matters" (keynote presentation at Open Web Application Security Project, April 4, 2012), accessed November 8, 2018, http://geer.tinho.net/geer.owasp.4iv12.txt.

39 Dan Geer, "You Are What You Eat," *IEEE Security & Privacy* 2 (July/August 2018), accessed November 8, 2018, http://geer.tinho.net/ieee/ieee.sp.geer.1807.pdf.

40 Peter Buneman, Sanjeev Khanna, and Wang-Chiew Tan, "Data Provenance: Some Basic Issues," accessed November 8, 2018, http://db.cis.upenn.edu/DL/fsttcs.pdf.

41 Full provenance and the tracking it implies can be used for both good and ill—another confirmation of the dual-use nature of all cybersecurity technology. See Dan Geer, "Provenance," *IEEE Security & Privacy* 88 (March/April 2016), accessed November 8, 2018, http://geer.tinho.net/ieee/ieee.sp.geer.1603.pdf; and Laney Salisbury and Aly Sujo, *Provenance: How a Con Man and a Forger Rewrote the History of Modern Art* (New York: Penguin, 2010), wherein data integrity, or the lack thereof, was taken to a stunning extreme.

42 In its most basic terms, diversity can be thought of as the range of storage systems, protocols, and tools that comprise an IT environment.

43 See Dan Geer, Rebecca Bace, Peter Gutmann, Perry Metzger, Charles P. Pfleeger, John S. Quarterman, and Bruce Schneier, "Cyber*In*security: The Cost of Monopoly," Computer and Communications Industry Association, September 2003, accessed November 8, 2018, http://geer.tinho.net/cyberinsecurity.pdf.

44 Thomas Dullien and Halvar Flake, "Security, Moore's Law, and the Anomaly of Cheap Complexity" (presentation at Conference on Cyber Conflict, May 2018, Tallinn, Estonia), accessed November 8, 2018, https://docs.google.com/presentation/d/17bKudNDduvN-7hWv7S84MiHUj2AnOPNbwjTM8euDC8w/mobilepresent?slide=3Did.g3b3c426762_0_41.

45 See, for example, Kathleen Fisher's work on quadcopter control. Kathleen Fisher, "Using Formal Methods to Enable More Secure Vehicles: DARPA's HACMS Program" (presentation at Tufts University, September 16, 2014), accessed November 8, 2018, http://verificationinstitute.org/wp-content/uploads/sites/28/2014/05/HACMS-Fisher.pdf.

46 See Sandy Clark's work on the inherent security downside of software reuse, for example: Sandy Clark, Stefan Frei, Matt Blaze, and Jonathan Smith, "Familiarity Breeds Contempt: The Honeymoon Effect and the Role of Legacy Code in Zero-day Vulnerabilities" (presentation at the 26th Annual Computer Security Applications Conference, Austin, TX, December 6–10, 2010), accessed November 8, 2018, https://dl.acm.org/citation.cfm?id=1920299.

47 Conversely, without adoption—by which we mean simply "use"—criticality cannot exist.

48 As just one illustration of this point, see Jeff Desjardins, "The Rising Speed of Technological Adoption," *Visual Capitalist,* February 14, 2018, accessed November 8, 2018, http://www.visualcapitalist.com/rising-speed-technological-adoption.

49 These models change in response to new data in production environments. And while, to our knowledge, such models are deployed in limited scenarios at present, we believe it is only a matter of time before they become widely adopted.

50 According to the European Centre for International Political Economy, in the decade preceding 2016, "The number of significant data localization measures in the world's large economies nearly tripled from 31 to 84." See Alan Beattie, "Data Protectionism: The Growing Menace to Global Business," *Financial Times,* May 13, 2018.

51 "China's Xi Jinping Says Internet Control Key to Stability," Reuters, April 21, 2018, accessed November 8, 2018, https://www.ndtv.com/world-news/chinas-xi-jinping-says-internet-control-key-to-stability-1840879. See also Beattie, "Data Protectionism": "China's Great Firewall has long blocked most foreign web applications, and a cyber security law passed in 2016 also imposed rules against exporting personal information, forcing companies including Apple and LinkedIn to hold information on Chinese users on local servers. Beijing has also given itself a variety of powers to block the export of 'important data' on grounds of reducing vaguely defined economic, scientific or technological risks to national security or the public interest."

52 See, for example, how a widespread reaction to the EU law has been to simply ban European IP addresses, described in Adam Satariano, "U.S. News Outlets Block European Readers over New Privacy Rules," *New York Times,* May 25, 2018, accessed November 8, 2018, https://www.nytimes.com/2018/05/25/business/media/europe-privacy-gdpr-us.html.

53 This is a point long predicted. See Jack Goldsmith and Tim Wu, *Who Controls the Internet? Illusions of a Borderless World* (New York: Oxford University Press, 2006).

54 The US Supreme Court's ruling in *South Dakota v. Wayfair Inc.,* 138 S. Ct. 2080 (2018), which paves the way for Internet sellers to pay sales tax based on the location of the buyer, not the seller, will only accelerate the trend of geocoding. See an overview of South Dakota v. Wayfair Inc., *SCOTUS* blog, accessed November 8, 2018, http://www.scotusblog.com/case-files/cases/south-dakota-v-wayfair-inc.

55 We are fully cognizant of the fact that many more unanswered questions remain.

56 David E. Sanger, *The Perfect Weapon: War, Sabotage, and Fear in the Cyber Age* (New York: Crown, 2018), 300.

57 The GDPR falls short in many ways, but it does get this aspect of data governance right. Specifically, Article 6 of the regulation outlines six—and only six—legal bases for data processing: GDPR Art. 6 § 1(a)–1(f) ("Lawfulness of Processing"). Other jurisdictions have taken note and are actively following the GDPR's lead.

58 See, for example, David M. Bitkower, deputy assistant attorney general, Criminal Division, Department of Justice, "Cyber Crime: Modernizing Our Legal Framework for the Information Age," testimony before the Senate Judiciary Committee's Subcommittee on Crime and Terrorism, July 8, 2015, accessed November 8, 2018, https://www.judiciary.senate.gov/imo/media/doc/07-08-15%20Bitkower%20Testimony.pdf. In language since repeated elsewhere in DOJ materials, Bitkower stated: "Current criminal law prohibits the creation of a botnet because it prohibits hacking into computers without authorization. It also prohibits the use of botnets to commit other crimes. But it is not similarly clear that the law prohibits the sale or renting of a botnet. In one case, for example, undercover officers discovered that a criminal was offering to sell a botnet consisting of thousands of victim computers. The officers accordingly 'bought' the botnet from the criminal and notified the victims that their computers were infected. The operation, however, did not result in a prosecutable U.S. offense because there was no evidence that the seller himself had created the botnet in question or used it for a different crime. While trafficking in botnets is sometimes chargeable under other subsections of the CFAA, this problem has resulted in, and will increasingly result in, the inability to prosecute individuals selling or renting access to many thousands of hacked computers."

59 That is not to say that this extension will be an afternoon's work. Disambiguating benign from malign intent will be no easier here than in other legal settings.

60 By "formalization" we mean both the ability of organizations to require identification *and* their ability to engage in the identification themselves (independently), both in the private sector and in the public sphere. Our growing inability to control the data we individually generate directly corresponds to our ability to be identified without our consent or even our knowledge.

61 See Reetika Khera, "Why India's Big Fix Is a Big Flub," *New York Times,* January 21, 2018, accessed November 8, 2018, https://www.nytimes.com/2018/01/21/opinion/india-aadhaar-biometric-id.html; and an overview of Estonia's "e-identity" program, accessed November 8, 2018, https://e-estonia.com/solutions/e-identity/id-card.

62 Which is exactly what some are already doing. Writing in *Wired,* for example, the executive director of the digital rights nonprofit organization Access Now warns of the dangers of digital IDs while stating flatly, "Digital IDs will become necessary to function in a connected digital world." Brett Solomon, "Digital IDs Are More Dangerous Than You Think," *Wired,* September 28, 2018, accessed November 8, 2018, https://www.wired.com/story/digital-ids-are-more-dangerous-than-you-think. See also Evgeny Morozov, "The Case for Publicly Enforced Online Rights," *Financial Times,* September 27, 2018.

63 See Hanna Hoikkala and Amanda Billner, "People in Sweden Now at Risk of Losing Access to Notes," *Bloomberg,* February 27, 2018, accessed November 8, 2018, https://www.bloomberg.com/news/articles/2018-02-28/swedes-now-at-risk-of-losing-access-to-cash-in-parts-of-country; David Crouch, "'Being Cash-free Puts Us at Risk of Attack': Swedes Turn Against Cashlessness," *Guardian,* April 3, 2018, accessed November 8, 2018, https://www.theguardian.com/world/2018/apr/03/being-cash-free-puts-us-at-risk-of-attack-swedes-turn-against-cashlessness; Maddy Savage, "The Swedes Rebelling Against a Cashless Society," *BBC,* April 6, 2018, accessed November 18, 2018, http://www.bbc.co.uk/news/business-43645676.

64 Jonathan Levin, "Cash Is in Short Supply in Storm-Ravaged Puerto Rico," *Bloomberg,* September 25, 2017, accessed November 8, 2018, https://www.bloomberg.com/news/articles/2017-09-25/king-cash-may-reign-for-weeks-in-storm-ravaged-puerto-rico; Frances Stead Sellers, Kevin Begos, and Katie Zezima, "After Hurricane Michael, Panama City Residents Cope With No Power, Cash-only Transactions and Baby-wipe Showers,"

Washington Post, October 21, 2018, accessed November 8, 2018, https://www.washingtonpost.com/national/hurricane-michael-left-many-in-panama-city-without-power-water-or-internet-putting-many-into-survival-mode/2018/10/21/7c54c39c-d316-11e8-8c22-fa2ef74bd6d6_story.html?utm_term=.002abc4e22c2.

65 See Dan Geer, "A Rubicon," *Lawfare* (blog), February 5, 2018, accessed November 8, 2018, https://lawfareblog.com/rubicon. Note that maintaining an analog option is both risk-reducing and a preserver of choice, the most fundamental of human rights.

66 The example of information security around elections is a case in point. Despite trends to the contrary, consensus suggests that all digital voting methods should have analog counterparts. See also Sen. Angus King (I-ME), Securing Energy Infrastructure Act, S. 79, 115th Cong. § 3 (2017), which calls for analog controllability of the electric grid, accessed November 8, 2018, https://www.congress.gov/bill/115th-congress/senate-bill/79; and Richard Danzig, "Surviving on a Diet of Poisoned Fruit: Reducing the National Security Risks of America's Cyber Dependencies," Center for a New American Security, July 2014, accessed November 8, 2018, https://www.cnas.org/publications/reports/surviving-on-a-diet-of-poisoned-fruit-reducing-the-national-security-risks-of-americas-cyber-dependencies: "Because cyber systems cannot assuredly be reliable, a strong presumption should be created that critical systems integrate non-cyber safeguards in their access and operation."

67 We readily acknowledge that this is a nontrivial and innately political decision. How would you impose a finite lifetime? A device abruptly fails? A user commits a crime by not throwing it away? What is the update requirement? How often? With what rigor? What if the vendor goes out of business (as has happened with public-key infrastructure certifying authorities)? How is the update channel to be protected from subversion? What does audit mean? Is this a research-grade question? Our key point is not that the answers to such questions are clear—or even currently exist—but simply that there must be laws around which such answers can be formed. This is far from the case today.

68 That theorem has been defined as follows: "It is only possible to simultaneously provide any two of the three following properties in distributed Web-based applications: *consistency* (C), *availability* (A), and *partition* tolerance (P)." Simon S. Y. Shim, "The CAP Theorem's Growing Impact," *IEEE Computer* 45, no. 2 (February 2012): 21, 22, accessed November 8, 2018, https://ieeexplore.ieee.org/document/6155651. See also Seth Gilbert and Nancy Lynch, "Brewer's Conjecture and the Feasibility of Consistent, Available, Partition-tolerant Web Services," *ACM SIGACT News* 33, no. 2 (June 2002): 51, accessed November 8, 2018, https://dl.acm.org/citation.cfm?id=564601; and Eric A. Brewer, "Towards Robust Distributed Systems" (keynote presentation at Principles of Distributed Computing Symposium, Portland, Oregon, July 2000), slides available at https://people.eecs.berkeley.edu/~brewer/cs262b-2004/PODC-keynote.pdf, accessed November 8, 2018.

69 Federal Trade Commission, "Mobile Security Updates: Understanding the Issues," February 2018, p. 2, accessed November 8, 20018, https://www.ftc.gov/system/files/documents/reports/mobile-security-updates-understanding-issues/mobile_security_updates_understanding_the_issues_publication_final.pdf.

70 This approach is, in some circles, referred to as "evidence-based policy making," which retains a strikingly bipartisan support base. See Commission on Evidence-Based Policymaking, "The Promise of Evidence-Based Policymaking," September 2017, accessed November 8, 2018, https://www.cep.gov/content/dam/cep/report/cep-final-report.pdf.

71 Cybersecurity Act of 2015, 6 U.S.C. §§ 1501–1533 (2015), accessed November 8, 2018, https://www.dni.gov/index.php/ic-legal-reference-book/cybersecurity-act-of-2015.

72 Reports indicate, for example, that only six private sector organizations share threat information with the government under CISA as of 2018. Joseph Marks, "Only 6 Non-Federal Groups Share Cyber Threat Info with Homeland Security," *Nextgov,* June 27, 2018, accessed November 8, 2018, https://www.nextgov.com/cybersecurity/2018/06/only-6-non-federal-groups-share-cyber-threat-info-homeland-security/149343.

73 See, for example, Office of Ron Wyden, US Senator for Oregon, "Wyden Pushes for Stronger Security in Collection of Personal Information," news release, May 15, 2017, accessed November 8, 2018, https://www

.wyden.senate.gov/news/press-releases/wyden-pushes-for-stronger-security-in-collection-of-personal-information.

74 We acknowledge the tension between this recommendation (to use the identity-obscuring capability of PETs) and our argument for embracing digital identification. More specifically, if we admit to—or, in the strongest sense, advocate for recognizing—the narrowing utility of anonymity, what is it that's left to preserve? Here, we distinguish between organizational data sharing, where the goal is generally to understand aggregate trends, and transactional activities between individuals and organizations in which anonymity is less and less valuable (or attainable). We refer to the former circumstance in this subsection, and the latter in the subsection entitled "embrace digital identification." More simply, when the anonymity of individuals is both valuable and *feasible,* PETs can and should be put to use. Otherwise, not.

75 Federal Aviation Administration, "Flying in Flat Light and White Out Conditions," 2001, accessed November 8, 2018, https://www.faasafety.gov/gslac/alc/libview_normal.aspx?id=6844.

The preferred citation for this publication is Andrew Burt and Daniel E. Geer, Jr., *Flat Light: Data Protection for the Disoriented, from Policy to Practice,* Hoover Working Group on National Security, Technology, and Law, Aegis Series Paper No. 1816 (November 20, 2018), available at https://www.lawfareblog.com/flat-light-data-protection-disoriented-policy-practice

About the Authors

ANDREW BURT

Andrew Burt is chief privacy officer and legal engineer at Immuta and a visiting fellow at Yale Law School's Information Society Project. Previously, he was special adviser for policy to the head of the FBI Cyber Division, where he served as lead author on the FBI's after-action report on the 2014 attack on Sony. He has published articles on technology, history, and law in the *New York Times,* the *Financial Times, Slate,* and the *Yale Journal of International Affairs*. He holds a JD from Yale Law School and a BA from McGill University.

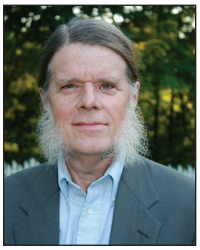

DANIEL E. GEER, JR.

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

Working Group on National Security, Technology, and Law

The Working Group on National Security, Technology, and Law brings together national and international specialists with broad interdisciplinary expertise to analyze how technology affects national security and national security law and how governments can use that technology to defend themselves, consistent with constitutional values and the rule of law.

The group focuses on a broad range of interests, from surveillance to counterterrorism to the dramatic impact that rapid technological change—digitalization, computerization, miniaturization, and automaticity—are having on national security and national security law. Topics include cybersecurity, the rise of drones and autonomous weapons systems, and the need for—and dangers of—state surveillance. The group's output will also be published on the Lawfare blog, which covers the merits of the underlying legal and policy debates of actions taken or contemplated to protect the nation and the nation's laws and legal institutions.

Jack Goldsmith and Benjamin Wittes are the cochairs of the National Security, Technology, and Law Working Group.

For more information about this Hoover Institution Working Group, visit us online at http://www.hoover.org/research-teams/national-security-technology-law-working-group.

Hoover Institution, Stanford University
434 Galvez Mall
Stanford, CA 94305-6003
650-723-1754

Hoover Institution in Washington
The Johnson Center
1399 New York Avenue NW, Suite 500
Washington, DC 20005
202-760-3200

HOOVER INSTITUTION