A Plea: The Case for Digital Environmentalism

Authors

Andrew Burt

Daniel E. Geer, Jr.

CSET CENTER *for* SECURITY *and* EMERGING TECHNOLOGY

November 2022

CSET occasionally provides a forum for guest writers to offer their expertise and opinions on emerging technology topics. They may describe a looming challenge or offer new recommendations regarding existing ones. The goal is to present fresh ideas, spark debate, and shape future research. Share your views on this paper with us at cset@georgetown.edu.

Executive Summary

The digital world is increasingly defined by unmitigable risks. Our data cannot be protected at scale. Systems make decisions that their programmers will never understand. It is no longer rare to encounter failures that cannot be addressed even after they are detected. With the fabric of human society so thoroughly digital—driving decisions in healthcare, finance, national security, transportation, and more—nothing less than long-term societal stability is at stake.

In this paper, we argue that digital environmentalism is the only answer to these woes. By “digital environmentalism,” we mean placing individual decisions within the broader context of the risks each generates and ensuring that these risks are distributed by intent rather than by happenstance.

In particular, we contest that the digital hazards we now face are a result of core assumptions about our world—key principles through which we view our basic relationship to technology. We call these assumptions grand “articles of faith,” and provide an overview of three such assumptions in Section 1.

The first article of faith is that there was, and continues to be, a meaningful distinction between the digital and the physical worlds. Put in simpler terms, the idea is that because it once made sense to separate the digital from the physical, it makes sense to *continue* to do so.

Second is the assumption that technology will always be able to ameliorate our woes—even when technology itself is the cause of our problems. Stated more broadly, this assumption has led us to a world in which technology solves one problem, only to create another, to which we apply a new technological fix, ad infinitum.

Last is the assumption that limitless growth can and should lead to limitless innovation—that is, innovation driving positive feedback loops within itself and without end.

Each of these articles of faith is, to put it simply, mistaken. Each in different ways, and each for different reasons. But each is contributing to a digital world that is fundamentally unsafe for those who populate it.

So what do we suggest? How, given the depth of our problems, are we to secure the world that we technologists are so busily creating?

We argue for three core solutions, knowing full well that much more action is needed. We do so in Section 2.

The first lies in empowering a single regulatory authority—be it the Federal Trade Commission or a new Digital Safety Administration—with the expertise and ability to enforce the limits in cyberspace that we believe are necessary. Such limits involve the creation and enforcement of standards for software development, and some form of mandatory auditing for portions of deployed code.

Second is equipping such an agency with tools to enforce core principles of sustainability. This means ensuring that new code and devices create a meaningful improvement in utility compared to existing alternatives, and are not simply developed for novelty's sake. It also means forcibly retiring unmaintained code when it is abandoned, as well as mandating accountability through independent assessments, or "red teams," that are focused on the risks of new technology *before* that technology is deployed.

Last is the removal—by Congress, the courts, regulatory agencies, or a combination—of software vendors' ability to disclaim liability in its entirety from their products. Holding software vendors to the same standards as every other industry will, over time, reduce the ability of flawed software to be introduced into the market.

Is it too late to secure the digital world? Is changing our collective approach to technology even feasible?

The momentum is against us; we readily admit that fact. The articles of faith we describe in this paper are deeply held and even more deeply ingrained. It will take significant commitment if we are to reverse these trends and to create a more sustainable digital world. We do not doubt any of these facts.

And yet so, too, are the stakes high. We now deploy code for almost every activity, and we are fast approaching a world where major portions of the software we rely on are incomprehensible and bug-ridden. A litany of failures therefore awaits us.

The time to take action, in other words, is now—before it is too late.

Introduction

We are, all of us, swimming in a sea of software and devices but with no perspective on where we are headed or the broader dangers we face.[1] It is neither late nor an exaggeration to say that the more we immerse ourselves in the digital world, the more severe the risks. As Elroy Dimson put it, "risk is that more things can happen than will."[2]

And here we are, madly increasing the number of things that can happen.

As a result, we now face a world full of unmitigable risks. Our data cannot be protected at scale. Systems make decisions that their programmers will never understand. It is no longer rare to encounter failures that cannot be addressed even once they are detected.[3] With the fabric of human society so thoroughly digital—driving decisions in healthcare, finance, national security, transportation, and more—nothing less than long-term societal stability is at stake.[4]

Digital environmentalism is, we contest, the answer—the only answer—to these woes. By "digital environmentalism" we simply mean placing individual decisions within the broader context of the risks each generates and ensuring that these risks are distributed by intent rather than by happenstance.

Every line of code that is deployed, every new IP address that is connected, every microchip that is printed, every self-modifying digital artifact buried inside something else—each carries with it implications for the digital environment. Evaluating each line of code, device, and decision in relation to their impact, and forcing their risks to be purposefully distributed, is the only way to prevent the failures we are otherwise sure to confront. Just as no one source of atmospheric carbon can change the climate, only the calculus of many can.

What Digital Environmentalism Is—And What It Is Not

When we use the term "environmentalism," we also mean *conservation*, in the sense of prioritizing the preservation of the environment, as well as *conservatism*, in the sense of encouraging gradual improvements that de-emphasize rapid change. Digital environmentalism is therefore, in our view, an approach to technological adoption that factors in both long-term (e.g., future users) and collective risk (e.g., societal impact) into how new technologies are developed and deployed. The question new technologies should elicit is therefore not simply "how useful is it?" but also "whom might it harm and when?"

Defining digital environmentalism by negation, however, may be a more practical way to make our point. In that sense, here is what digital environmentalism is not: it is not developing new technologies and rushing them to market because they are novel; it is not assessing technology for risk after it is deployed, and only then applying patches or other mitigants; it is not praising innovation for innovation's sake. Digital environmentalism is therefore admittedly (and fundamentally) at odds with our current approach to technological adoption—an approach that places new and fast over sustainable, jeopardizing our privacy, security, and even our safety in the process.

For many, the term "digital environmentalism" might sound odd—a combination of two terms that are not typically associated with each other. But these terms *should* be, and the oddity of their association is precisely our point, not least because the digital world impacts our environment in many ways. For starters, the problem we face is far from abstract. Carbon emissions from internet-connected entities account for roughly 3 percent of all global electricity usage and 2 percent of greenhouse emissions.[5] The digital world itself is built on silicon and relies on precious metals that have limited supply. An iPhone purportedly contains 75 elements from the periodic table, roughly two thirds of the 118 elements known to science.[6]

And then there is cyberspace itself: a combination of networks, applications, and sensors that digitally capture ever more data about ever more activities.[7] This environment requires both safeguarding and protection and yet receives, when it comes down to it, precious little of each—an assertion that is by now inarguable, based simply on the prevalence of successful cyberattacks.[8] It is no stretch to say that digital technology, the defining innovation of the last half century, has deep and unaddressed insecurities at its core.[9]

The questions, then, are twofold: How did we come to live in an environment so plagued by downsides? And, more immediately, what, if anything, should we do to address them?

We spend Section 1 diagnosing these problems, and Section 2 describing how and why we believe digital environmentalism, in its broadest sense, is their only cure. For readers who are less interested in our diagnosis than in our proposed cure, you will be forgiven for skipping directly to Section 2.

1. Hidden Assumptions, Hidden Risks

If we find ourselves living in a world we cannot protect, it is not because we—any of us—wanted to find ourselves here. Environmental hazards are, by their nature, collective threats; they pose profound dangers to us all. They are, in effect, integral calculus—the infinite sum of infinitesimals.

These hazards, including those in the digital sphere, are a result of assumptions about our world—that this or that individual action or event is too small to matter. Those assumptions have led us to the present circumstance.

Throughout this essay we call these assumptions grand “articles of faith.”[10] Left unexamined, these articles of faith are leading us to a world in which we face an insurmountable number of unmitigable risks. There are many such articles of faith; we start with three.

Digital Versus Physical Environment

The first article of faith is the notion that the distinction between the digital and physical environments is both real and meaningful. Put in simpler terms, the idea is that because it once made sense to separate the digital from the physical, it makes sense to *continue* to do so.

The more digital technologies we adopt, the more of our actions are digitally captured and transmitted, meaning that, at minimum, it has become inconceivable to “opt out” of the digital world. At maximum, this means that our activities are as much digital as they are physical, if not more so.

A few examples:

Researchers have demonstrated the ability to detect heart and breathing rates through wireless internet signals, meaning that anyone within range of a wireless router is, theoretically, exposing their vital signs to surveillance.[11]

Or take smartphones, which are now a necessary part of life for many of us.[12] The simple act of keeping a cell phone powered-on over time creates a real-time map of the owner's activities, which can be tracked by third parties to determine who they are and what they are doing (both online and off).[13]

Indeed, even the activities we take that are not digitized—such as *refusing* to join a specific social media network—can reveal unintended, personal information that most would not intuitively think of as digital. Researchers have proven, for example, that they can infer who an individuals' friends, family, and professional contacts are based on the social media activities of those around them—even when an individual explicitly refrains from taking part in any one specific platform.[14]

> *"As digital activities become more important, and as data collection becomes increasingly pervasive, the once-bright line between digital and physical no longer exists."*

All of which, in turn, means that there is no such thing as the purely "digital" world in practice. As digital activities become more important, and as data collection becomes increasingly pervasive, the once-bright line between digital and physical no longer exists.[15] It is more and more common, for example, for cyberattacks themselves to create direct physical repercussions, as when a ransomware attack on a hospital allegedly led to a patient's death,[16] or when the NotPetya wiper collapsed global supply chains[17]—the list is long and growing fast.

It no longer makes any sense, in other words, to think that simply because the dangers of software are embedded in code, that such dangers are somehow less consequential. To imagine that we can still separate the digital from the physical is to ignore the most inconvenient truth.

Malthus's Grand Lesson

Over two centuries ago, an economist and cleric named John Malthus predicted that the human race would soon grow to a size that was unsustainable—that we humans would cease to be able to feed ourselves given existing environmental constraints and associated crop yields. Using deeply convincing arguments, Malthus showed how the human population grows exponentially over time while environmental resources grow linearly. In other words, huge portions of humanity were destined for malnourishment, famine, and ultimate extinction.

What Malthus did not know was that technology would prevent his apocalyptic predictions from coming true: the steam engine would make supply chains more efficient, enabling faster movement of goods across wider areas; developments in agricultural techniques would increase food yields without requiring more land; eventually, fertility technologies themselves, such as contraceptives, would enable more conscientious family planning for those who could access them.

The fact that Malthus was wrong—indeed, horrendously wrong—was largely due to technologies whose development he was unaware of and whose success he could not foresee. That dynamic has given rise to the second, central article of faith for most inhabitants of the industrialized world: that there is no problem that technology cannot fix. No matter how big the problem, the right technology—humanity's ability to collectively fashion objects into tools—will solve it, given enough time and the proper incentives. Mother Necessity will reliably birth inventions as required. Or so such thinking goes.

This approach to technology is on display everywhere. Bill Gates, the founder of Microsoft, has stated that the only way to combat climate change lies in developing entirely new "green" technologies.[18] It underlies the now-commonplace claims that artificial intelligence will solve humanity's biggest problems, from cybersecurity to disease. Indeed, it forms the foundation for modern economic theory, which assumes constant growth through limitless innovation—or, in other words, an infinite sequence of new technologies being invented to meet our needs.[19]

As Ronald Reagan once expressed it: "There are no great limits to growth because there are no limits of human intelligence, imagination, and wonder."

To which we raise the obvious question: What if these assumptions are wrong?

What if there is a point—a breaking point—where the problems we face cannot be solved by new technologies? And, furthermore, if this is a possibility, how would we know when we were nearing that point of inflection?

Stated more broadly, this assumption has led us to a world in which technology solves one problem, only to create another, to which we apply a new technological fix, ad infinitum. If technology broke it, in other words, then more technology must be able to fix it, whatever it was that was broken and whatever the resources required.[20]

There are, however, some problems that simply cannot be solved by new tools or technologies.[21] To readers who might object to this truth, we have only one response: welcome to reality. You may not like its limitations or its dictates, but your feelings are beside the point.

Limitless Growth, Limitless Innovation

An expectation of limitless growth assumes limitless innovation—that is, innovation driving positive feedback loops within itself. Should your desire be limitless growth then debating limits to growth is futile; limits would only harm the overall aim of growth.

From this standpoint, one might say that we face a binary choice between free-running, limitless innovation on the one hand and enforced stasis on the other.[22] Any constraint placed on innovation will toy with those positive feedback loops. We will never know what might have been had we not held back.

Let's start with some examples of this approach to illustrate our point. Here is how Jeff Bezos, among the most influential voices in the technology industry, views the problem of limitless innovation:

> Earth is finite, and energy usage has grown so much that it will soon . . . strain the resources of our small planet. That will leave us with a choice: accept static growth for humanity or explore and expand to places beyond Earth.[23]

Bezos' argument posits a binary choice: embrace limitless growth and, therefore, the need to conquer the universe, or constrain growth and thereby conserve our environment. (If the former sounds like the outlook of an alien species bent on consuming the universe's natural resources without end, everything in life is a matter of perspective.[24])

It is worth, of course, stating that this outlook is about much more than space travel—yes, it touches on traditional environmental concerns, but also deeply technological ones as well. The idea of a "carrying capacity" is relevant not only to humankind's extractive interaction with a finite planet, but is also a general inference about how much perturbation an ecosystem can endure without an extinction event.

Indeed, there is a direct link between the economic assumptions demanding endless growth and the specific technologies that shape the digital world. One of the main economic models shaping venture capitalists' approach to new technologies, for example, is explicitly focused on the ability of new inventions to be monetized and pumped into an ever-expanding market.[25] The very idea of an "information economy" was inspired by this approach to new technologies.[26]

The same assumptions famously hold true for physical hardware, not just software. Moore's Law long defined the computing industry's expectations for transistors, that the density of transistors on a microchip doubles every two years. This has, more or less, held true for the last five decades, and is based on the assumption that growth—and in this case, microchip innovation—can and will continue.[27]

It may seem obvious, but these assumptions drive everything from the $200 billion invested in early-stage startups last year, to the 932 billion microchips manufactured in 2020.[28] Without endless growth, there would be a limit—at least in theory—to the production of both digital and physical goods.

Indeed, change through innovation has become not just the mantra of the dissatisfied, but the core source of the adhesion between investors and technologists. Fast change is what births first mover advantage. Fast change promises to level the playing field between incumbents and upstarts.[29] Fast change is paranoia-inducing.[30] Fast change makes short-term thinking the only way to fly.[31]

The technology treadmill is perpetual, and we are to keep running on it, all of us, lest we fall over. Cue the Red Queen.[32]

And now we focus on risk: attempting to make safe, or even understand, a code base, a set of hardware components, or *any* environment is made difficult in proportion to its rate of change.[33] If that environment is constantly and rapidly expanding, its complexity and interdependencies ever increasing, then managing such risks becomes asymptotically impossible—rebuilding the landing gear while in flight, so to speak.[34] Sure, we can make progress around the edges, but the game of risk management is lost by definition before it even begins.

"By now, it should be clear that the problem is environmental in that term's broadest sense."

By now, it should be clear that the problem is *environmental* in that term's broadest sense: we cannot protect a set of resources we do not understand, whose complexity surprises even the most well-resourced actors, whose interdependencies are discovered only when deployed, and whose individual criticality is due not to design but rather to the simple calculus of mass adoption. We are, in a very profound sense, unsafe for as long as we live in an environment that we cannot quantifiably secure. And if we are unsafe now, our insecurity is destined to grow with each increment of complexity.[35]

We are addicts, all of us, adopting ever more technologies that cannot save us from the dangers they create. It is our articles of faith, each one of them, that have left us in an environment we cannot control, little less hope to secure.

2. What to Do?

Are there alternatives to limitless growth? Does the binary choice between stasis and innovation preclude a third way?

We have deep, it seems, and serious problems in the world of technology. For readers who have made it this far in this essay, that much is clear. The question, then, is how the assumptions we describe above combine to create the risks that are jeopardizing our environment. More to the point, what are we to do about them once discovered?

But first, two disclaimers: To begin with, the recommendations we make in this section are national in flavor—while all environmental threats are global, all politics are, as they say, local, which forms the reasoning behind our focus on the United States. We are under no illusion, however, about the scale and scope of the problems we describe. For that reason, it is worth emphasizing that the solutions to these issues must be global in the fullest sense.

Just as the global environment does not distinguish between jurisdictional boundaries, digital threats cut across, and in many ways transcend, national boundaries as well.

Second, the problems we describe—our assumptions about limitless growth, our increasingly mistaken distinction between physical and digital, our unwavering belief in the power of new technologies to ameliorate our woes—cannot be solved by focusing on digital problems alone. These same issues are also deeply intertwined with broader assumptions about economic growth and societal development.[36] Indeed, the idea of limitless innovation fueled by limitless resource extraction, leaving long-term environmental externalities unaddressed, is one of the organizing principles of modern society. If we are to truly course correct, then, these broader issues must also be confronted.

That said, however, we leave the larger work of questioning these cultural assumptions to others—and not just for convenience's sake. We focus here in this paper, and in the remaining sections, on what we know best: digital technology and its shortcomings.[37]

What Really Ails Us?

At this point, it is worth reminding readers that we are both technologists. Technology can be, and very often is, a *good thing*—improving the way we communicate, access and process information and the like. There is much to be optimistic about; this optimism explains why the two of us have dedicated much of our lives to digital technologies.

The problem has less to do with digital technologies themselves than it does with their relation to the broader environment and the articles of faith, outlined above, which shape that environment. Indeed, the fundamental issue at hand is how the assumptions we described in Section 1 distribute risks across actors, applications, and devices.

Digital technology is cheap to create and costly to clean up. This leads us to a world in which a new device is always worth more than a refurbished one; where new, backward incompatible software is incentivized over improving existing code; where AI, churning out endless predictions from a growing body of data, is viewed as inherently more valuable than static logic even if applied to the same problems because it is *newer* tech. If technology is good and limitless growth is the goal, the new will trump the old every time.[38]

The long-term implications to the environment do not stand a chance against the short-term worship of the new. Engineers have long had this dictum: "Fast, cheap, reliable—choose two."

A recasting of this dictum seems now to be in order.

Indeed, the risks of ever more and ever newer are all-too-frequently dispersed to those outside of any short-term transaction. It is not the software maker or the software purchaser who need to worry about end-of-life issues for much of the code that is used. By the time a problem emerges, the maker and the buyer will have moved on to newer things, selling, for example, their application to a bigger company (in the case of the vendor) or disposing of or giving away the device itself (in the case of the consumer).

Instead, vulnerabilities in the code or artifact become absorbed by the broader digital ecosystem, like microplastics washing into the ocean. The "environment" is expected to deal with the externalities of the transaction at no cost to the immediate parties. This is the moral hazard of the digital world, explaining why a huge percentage of data breaches are caused by third party vulnerabilities.[39]

This type of interdependence forms the root of systemic risks. When interdependence rises, risks become transitive though benefits do not—a common dependency does not inherently create a common benefit.[40] The positive feedback character of digital innovation causes those on the leading edge of adoption to have different risks from those who are laggards. Attacks go where exploitable targets concentrate; one or two versions behind "current" is probably the most dangerous place to be for widely used code.

A digital world marked by both growing interdependence and quickening change cannot be predictable enough to be understood, nor understood enough to be predictable.[41]

Fixing Our Problems

If it seems we are arguing that the sky is falling, we are not. If it seems we are arguing that our problems are so profound, and their causes so deep, that there are no easy fixes, we are. But there are fixes we can adopt nonetheless.

> *"Put simply, we must embrace a practical form of environmentalism more widely and apply its tenets to the digital world."*

What, then, do we advocate in the face of such dangers?

Put simply, we must embrace a practical form of environmentalism more widely and apply its tenets to the digital world.[42] At its core, this means applying *limits*—limits to how we engage with the environment and ensuring that risks are distributed across space and time in a way that can be said to be both intentional and consistent with the values we claim to hold. Innovation without limits cannot, by its very nature, be sustainable, meaning that growth without end is itself an unmitigable risk.

So what, in practice, might these limits look like? There are many forms of environmentalism that might be applied to the digital world. We outline three core tenets of one approach, knowing full well that our recommendations are merely an outline of what is needed.

Enter the Digital Safety Administration?

If placing limitations on the digital world is, as we argue, the most important step in reasserting control over our environment, then there must be an entity capable of enforcing such limits to begin with.

Dispersed enforcement, of the kind we are all too familiar with in the United States, will not suffice. Software systems are simply too pervasive to put our stock in a piecemeal approach where one entity regulates a specific sector, and another regulates another. It is foolish, therefore, to divide the regulation of software between the Food and Drug Administration, the National Highway Traffic Safety Administration, the Federal Deposit Insurance Corporation, et cetera, et cetera, all the while expecting a coherent outcome.

If there is to be efficacious software regulation, expertise and resources must be concentrated in a single agency or actor, one capable of supervising, standard setting, and enforcing regulatory actions against purveyors of digital risk.[43] As to whether that agency is new or old, we have no preference. The Federal Trade Commission has a long history of supervising new technologies; clarifying its enforcement powers over software writ large, rather than simply unfair and deceptive practices in general, might do the trick. Neither would we object to a new agency—perhaps entitled the Digital Safety Administration—with sole authority for supervising the creation, maintenance, and decommissioning of code.

This entity would have a long list of critical tasks. High on that list must be public audits of code, which would encompass the review and testing of portions of software systems to ensure that they are free from detectable errors or vulnerabilities, just as foodstuffs are audited for the presence of trace poisons. Audits of this sort might even parallel financial audits, where a provider generates its own audits and asserts its bill of health to the regulator in a binding fashion. There are many ways to engage in such audits; the methods matter far less than the fact that audits are mandated and occur in practice.[44]

Here we point to a key limitation: the digital world is far too large for direct supervision of every line of software that is deployed. That is true today, and the codebase of cyberspace will only grow over time—making this problem even more acute in the future. We are clear eyed about the obstacles confronting this type of supervision.

But these limitations need not impede the type of oversight that we are calling for. Regulatory agencies in the United States and elsewhere have long operated within such asymmetrical enforcement environments, where regulatory resources are significantly limited in comparison to the market each agency supervises. The Internal Revenue Service, for example, need not audit *every* taxpayer filing to ensure that most entities pay their taxes on time and in reasonably good faith; it need only audit *enough* taxpayers that the chances of an audit for each individual are not entirely negligible.

The same should apply to the software overseen by such an agency. It need not even audit entire codebases; a carefully selected sample set of code would do. Where errors or negligent development practices are detected, penalties should ensue and with enough regularity that no one entity or program can confidently ignore the consequences of violations.[45] And just as the IRS depends on the factual record-keeping of those it regulates, software providers could have factual record-keeping duties.[46]

When it comes to digital versus physical environmentalism, it is also worth stressing that digital environmentalism has upsides that physical environmentalism does not—making a Digital Safety Administration easier to implement in practice than its analogs in the physical world.[47] In the digital world, for example, risk distribution is a much more intuitive task precisely because cyberspace is by itself an artificial construct. Every piece of the digital fabric, so to speak, was created by a human in some particular time and at some particular place. That human can be held responsible. Assigning an agency responsibility over risk distribution, and enforcing its mandate, is therefore a more straightforward task for digital risks.[48]

Three Principles for Sustainability in Cyberspace

Beyond audits of code, the enforcement of limitations on cyberspace must also be premised on clear principles of sustainability, by which we mean a few things. First is ensuring incentives for reuse. If we are to protect an environment defined by its finitude, we must, as a society, be able to *justify* the creation and use of new materials. If a device or application is new, its utility should therefore be a dramatic improvement on what preceded it. We cannot, in other words, adopt the new for the sake of novelty, generating unnecessary risks and waste without sufficient collective benefit.

Instead, we must be intentional about the utility of the materials that we use and the tools we adopt in relation to their risks. This runs counter to the all-too-common development process, where security is added on if and only if the software system finds a market sufficient to justify the cost of security. As with other matters of regulation, the impact of market concentration will itself be part of the regulator's oversight.[49]

This means that, as a rule, developers should be incentivized to update existing code rather than to start from scratch. New libraries should be separated from old ones; longevity should not simply be a virtue in theory, but a benefit in fact. It should therefore cost less to reuse something that already exists than to create something new.[50]

There are a number of ways to achieve these goals, each with distinct downsides but each, also, attempting to serve the same ends: there could be some type of tax on new materials, levied in proportion to the risks such materials generate; there could be reduced liability for items that incorporate existing materials, such that if an audit of the type we described above were to occur, the penalties for any violations would be reduced; there could even be public funding devoted to sustaining the types of materials that already get used over and over again, such as open-source libraries or the refurbishment of old devices.[51] This list is far from comprehensive. Our point is that we have many options, but have not come close to attempting—little less exhausting—a fraction of them.[52]

Second is mandating limits on the lifetime of unmaintained code. Longevity is indeed a benefit, but only if that longevity is not synonymous with neglect. All too often, old code is bug-ridden *when forsaken*, meaning that its very use poses risks to the broader environment it is deployed within. Just as the world is awash in trash that does not biodegrade, the digital world is awash in artifacts whose neglect renders them both immortal and unfixable.

If, in fact, digital ephemera are more like milk than wine,[53] as we contend—needing a kind of refrigeration and enforceable use-by dates—there are a host of requirements we might implement to mitigate their dangers: placing limits on the lifetime of unmaintained code and hardware, for starters, but also mandating auto-updates (the digital world's equivalent to product recalls), ongoing penetration testing, well curated and verifiable deployment inventories, software bills of materials, testable reachability guarantees for remote management interfaces, and more.

Above all, however, should be the requirement that abandoned code must be subject to legally binding forfeiture processes.[54] Fines, for example, should be levied for software that is left unpatched but deployed in live environments. The longer the neglect, the steeper the penalty should be. As we have asserted elsewhere, if a device has an IP address, it must either accept updates or have a finite lifetime.[55] Retire or repair unmaintained code—there must not be a third option.[56]

We note that there is no contradiction between our argument for the conservation of existing code and for the forcible retirement of code that is no longer well-maintained. Both requirements can, in fact, be met, simply by ensuring that the code that is deployed is kept serviceable, and that new code only be deployed when justifiable. We should strive to sustain both the code we deploy and the devices software lives on.[57]

Third is the notion long relied upon by the information security community, which is the use of red teams—or using third parties to identify vulnerabilities and risks that would otherwise be left unaddressed. Risk, after all, is a matter of perspective. It is hard to have the right perspective when focusing on speeding up the time to market, just as it is hard to assess a product for shortcomings when your promotion depends on its success. Some call this accountability; we call it a form of necessary quality control.

This means that entities who are not responsible for commercial success must be required to assess software for risks. Such entities may be external third parties; they may also be independent developers within the same organization that developed the code. The regulatory framework governing the use of algorithms in finance even has a specific word for this type of review: *effective challenge*.[58] The Federal Trade Commission already encourages these types of accountability mechanisms.[59] Whatever we call it, the idea is not new. The key lies in incentives: ensuring that software is reviewed for errors by reasonably independent parties, before it is deployed, will serve as a bulwark against the types of errors that can lead to catastrophic risks.

And here we add one note: by software, we mean *all* software. If it is connected to other devices or systems, that code is part of the digital environment and therefore poses collective risks. Just like littering in public, fines must accompany harmful behavior so as to incentivize the good.

The Buck Stops . . . Somewhere

The current, entirely pervasive use of disclaimers in end user license agreements must end. It is all too common in the digital world to waive nearly all liability when selling or distributing code—a fact that is beneficial for its purveyors but harmful for everyone else.[60]

Here, for example, is how WhatsApp, one of the most used applications on the planet, describes the guarantees it provides for its code:

> We are providing our services on an “as is” basis without any express or implied warranties, including, but not limited to, warranties of merchantability, fitness for a particular purpose, title, non-infringement, and freedom from computer virus or

> other harmful code. We do not warrant that any information provided by us is accurate, complete, or useful, that our services will be operational, error free, secure, or safe, or that our services will function without disruptions, delays, or imperfections.[61]

This language is hardly unique to WhatsApp or its parent company Meta. Anywhere someone is using software, this type of waiver is likely in effect, meaning that the creators of the digital world are largely unaccountable for the harms their activities create.[62]

What other service or product is so impactful and yet so fancy free?

The answer is none. Nearly every other industry is held to higher standards and therefore remains more accountable for their activities.[63] But because the digital world is viewed, mistakenly, as inherently different from the physical world, and because incentive structures reward fast development of new code, digital technologies receive a free pass. How exactly this revocation should occur in practice, we leave to others. Much like the history of liability in the automobile industry, there are myriad ways to limit these types of disclaimers, from regulatory activity at the federal, state, or local levels to increased scrutiny by the courts.[64]

If we are to protect the digital world, this free pass must be revoked. For this reason, there must be meaningful limitations on what liabilities software developers can waive, and in what circumstances.[65]

Technology's Environmental Paradox

The momentum is against us; we readily admit that fact. The assumptions we described in Section 1 are deeply held and even more deeply ingrained. It will take significant commitment if we are to reverse these trends and to create a more sustainable digital world.

"We are fast approaching a world where major portions of the software we rely on are incomprehensible and bug-ridden."

And yet so, too, are the stakes high. We now deploy code for almost every activity, and we are fast approaching a world where major portions of the software we rely on are incomprehensible and bug-ridden.

A litany of failures therefore awaits us—exploited by nation states, criminal enterprises, or simply arising by happenstance. This is a world full of failures we cannot afford, undermining the trust and reliability of the systems we use in finance, healthcare, and beyond, and jeopardizing the environment in which we now live.

But these types of failures cannot remain inherent if we are to maintain a baseline sense of security in both the digital and physical worlds. This means that we are only secure to the extent that we can respond to failures when they arise. Silent failures, on the other hand, create vulnerabilities we cannot, by their very nature, address.

All of which leads us to a fundamental paradox: the only way to protect the technology we create is to protect the environment in which it is deployed. And the only way to protect that environment is to place limits on the creation and use of that very same technology.

Meaning that we, as digital environmentalists, are on the same side as all of those who argue in favor of limitless innovation. Only by prioritizing the long-term sustainability of our environment can we truly reap the short-term benefits of the tools we create. Technology and the environment cannot be disentangled, in other words, for the same reason that the digital is now the physical for all practical purposes.

Sustainability must not therefore be a catchphrase, or a marketing ploy, but a practical method to ensure that innovation does not lead us to catastrophic risks. Otherwise, we jeopardize both the environment we live in and the very technology we prize.

Authors

Andrew Burt is co-founder and managing partner of BNH.AI, a boutique law firm focused on artificial intelligence and analytics, as well as a visiting fellow at Yale Law School's Information Society Project. Previously, he served as chief legal officer at Immuta and as special advisor for policy to the head of the FBI Cyber Division. He holds a JD from Yale Law School.

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 a senior fellow at In-Q-Tel. His collected writings can be found at http://geer.tinho.net/pubs.

Acknowledgments

For feedback and assistance, we would like to thank, among others, Jack Balkin, John Bansemer, Michael Daniel, Oona Hathaway, Drew Lohn, Igor Mikolic-Torreira, Joanna Murdick, and Matthew Schaffer.

Document Identifier: doi: 10.51593/2022CA005

Endnotes

[1] Andrew Burt and Daniel E. Geer, Jr., “Flat Light: Data Protection for the Disoriented, From Policy to Practice,” Aegis Series Paper No. 1816, Hoover Institution Essay, November 2018, https://www.hoover.org/sites/default/files/research/docs/burtgeer_flatlight_revisednov20_webreadypdf_final.pdf.

[2] Elroy Dimson, quoted in Peter Bernstein, *Against the Gods: The Remarkable Story of Risk* (Hoboken: Wiley, 1996). The things occurring, of course, being those with negative consequences rather than beneficial results.

[3] For this reason, even determinism itself seems quaint.

[4] If we fail, it will be for the most inexorable of reasons—declining marginal utility of increased complexity. See Joseph Tainter, *The Collapse of Complex Societies* (Cambridge: Cambridge University Press, 1990).

[5] See “The Carbon Footprint of the Internet,” Climate Impact Partners, published April 22, 2021, https://www.climatecare.org/resources/news/infographic-carbon-footprint-internet.

[6] Carolyn Wilke, “Materials of the last century shaped modern life, but at a price,” *Science News*, January 28, 2022, https://www.sciencenews.org/article/materials-science-history-modern-daily-life-environment; (“‘The iPhone contains about 75 elements from the periodic table—a huge proportion of all the atoms that we know about in the universe are in an iPhone,’ says Ploszajski, the materials scientist. Some of those are rare-earth elements, a set of 17 metallic elements mostly on the outskirts of the periodic table. Though they are difficult to mine and process, rare earths are sought after because they lend unusual magnetic, fluorescent and electrical properties to materials made from them. Neodymium, for example, mixed with other metals makes the strongest magnets known. These magnets make your cell phone vibrate and its speakers produce sound.”) Note, however, that our focus throughout this paper will be primarily on digital technologies and the digital world—notwithstanding the increasingly meaningless distinction between digital and physical, as we argue in Section 1.

[7] Indeed, cyberspace itself is increasingly inseverable, in the literal sense, from the physical world. More and more medical devices are internet connected, for example, as are the industrial systems that connect the electric grid, distribute water, and operate sewage facilities. If the device or system was created in the last decade, chances are it is connected directly or indirectly to the digital world. For further discussion of the impact of interconnection, see Daniel E. Geer, Jr., “A Rubicon,” Aegis Series Paper No. 1801, Hoover Institution Essay, February 2018, https://www.hoover.org/research/rubicon.

[8] The list of references detailing the privacy and security vulnerabilities in cyberspace is endless. For starters, see Andrew Burt, “Nowhere to Hide: Data, Cyberspace, and the Dangers of the Digital World,” Digital Future Whitepaper Series, Yale Information Society Project, December 2020, https://law.yale.edu/sites/default/files/area/center/isp/documents/white_paper_2020_nowhere_to_hide_burt_yls_isp_digital_future.pdf.

[9] This is what is meant by the term of art "security debt," a fact unignorable if for no other reason than that a bug is exercised by an accident while a vulnerability is exercised by an enemy. See Simo Huopio, "A Quest for Indicators of Security Debt" (and references therein), *The Cyber Defense Review* 5, no. 1 (Spring 2020), https://cyberdefensereview.army.mil/Portals/6/CDR%20V5N1%20-%20FULL_WEB_1.pdf?ver=2020-03-23-113809-503.

[10] While such articles of faith have shaped the modern, industrialized approach to technology, we do not mean to suggest that they are universally held by *everyone*. In other words, while not everyone "shares the faith" in such articles, our world (and especially the digital environment) is deeply impacted—if not defined—by them.

[11] Yu Gu et al., "WiFi-based Real-time Breathing and Heart Rate Monitoring during Sleep," arXiv preprint arXiv:1908.05108 (2019), https://arxiv.org/abs/1908.05108. Given the ubiquity of wireless internet, there are not many places where individuals are not so exposed. The assumption being, of course, that the types of invasive surveillance technologies that are being proven in a lab setting will become more pervasive in the future. See, for example, the discussion of gait analysis, authorship identification, and other privacy invasive technologies in Burt, "Nowhere to Hide."

[12] Eighty-five percent of Americans, for example, now own a smartphone. "Mobile Fact Sheet," Pew Research Center, April 7, 2021, https://www.pewresearch.org/internet/fact-sheet/mobile.

[13] Examples of such data violating individual privacy are commonplace; one notable example includes the outing of a prominent Catholic priest for using the dating app Grindr, based on nothing other than the location data of his phone and information that journalists were able to purchase from third-party tracking companies. See Molly Olmstead, "A Prominent Priest Was Outed for Using Grindr. Experts Say It's a Warning Sign," *Slate*, July 21, 2021, https://slate.com/technology/2021/07/catholic-priest-grindr-data-privacy.html. Such information can even be used to determine individuals' gaits. See Nantakrit Yodpijit et al., "The use of smartphone for gait analysis," in *3rd International Conference on Control, Automation and Robotics (ICCAR)* (2017), https://ieeexplore.ieee.org/document/7942756.

[14] We are, in other words, all intelligence officers now. James P. Bagrow, Xipei Liu, and Lewis Mitchell, "Information flow reveals prediction limits in online social activity," *Nature Human Behaviour* 3 (January 2019), https://www.nature.com/articles/s41562-018-0510-5; ("Our results have distinct privacy implications: information is so strongly embedded in a social network that, in principle, one can profile an individual from their available social ties even when the individual forgoes the platform completely."). See also Dan Geer, "We Are All Intelligence Officers Now," (lecture, RSA Conference, San Francisco, February 2014).

[15] Shoshana Zuboff, *The Age of Surveillance Capitalism* (London: Profile Books, 2019).

[16] Kevin Poulsen et al., "The Case of the First Alleged Ransomware Death," *The Wall Street Journal*, September 30, 2021, https://www.wsj.com/articles/ransomware-hackers-hospital-first-alleged-death-11633008116.

[17] Andy Greenberg, "The Untold Story of NotPetya," *Wired*, August 22, 2018, https://www.wired.com/story/notpetya-cyberattack-ukraine-russia-code-crashed-the-world.

[18] As Gates argued last year in the *Financial Times*, "innovation is the only way the world can cut net greenhouse gas emissions from roughly 51bn tonnes per year to zero by 2050," (Bill Gates, "Funding clean technology is the way to avoid climate disaster," *Financial Times*, October 31, 2021, https://www.ft.com/content/ea71f4f8-e5d8-4324-a42c-8fa09ccb1cc5).

[19] An assumption, we note, upon which both authors' retirement savings may be said to depend.

[20] As stated by a quote often attributed to computer scientist David Wheeler: "Any problem in computer science can be solved with another level of indirection, except for the problem of too many layers of indirection."

[21] Just as some questions are known to be undecidable, some problems are known to be unsolvable, meaning that no technology can itself solve them. See, for example, Gödel's incompleteness theorems. For a more general overview of so-called undecidability problems, see Toby S. Cubitt, David Pérez-García, and Michael Wolf, "The Unsolvable Problem," *Scientific American*, October 1, 2018, https://www.scientificamerican.com/article/the-unsolvable-problem/.

[22] In reality, the choice is starker than that—because our current approach to innovation is fundamentally unsustainable, so-called limitless innovation is more akin to (eventual) environmental collapse. Is it possible, on the other hand, to have "clean" growth and therefore sustainable innovation? Yes, but only with clear limits—limits being the central requirement of our proposed approach to digital environmentalism, as we argue below.

[23] As described by Jobs' biographer, Walter Isaacson, *Invent and Wander* (Boston: Harvard Business Review Press, November 2020), 19.

[24] This approach to natural resources is how Thomas Whitmore, the fictional president in the movie *Independence Day*, describes the aliens invading Earth: "They're like locusts. They travel from planet to planet, their whole civilization. After they've consumed every natural resource, they move on. And we're next." See http://www.dailyscript.com/scripts/id4.html.

[25] The theory is called "New Growth Theory," the progenitor of which was awarded the Nobel Prize in 2018. See Paul M. Romer, "Endogenous Technological Change," *The Journal of Political Economy* 98, no. 5 (October 1990), https://web.stanford.edu/~klenow/Romer_1990.pdf. We note, however, that within the world of economics, a countermovement to such widespread thinking has also emerged, perhaps most notably with the 1972 report "Limits to Growth" by the Club of Rome.

[26] Andrew J. Sutter, "Unlimited Growth and Innovation: Paradise or Paradox?," SSRN Paper, November 2010, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1709285.

[27] While there have been mumblings about the end of Moore's law, it appears that the technology industry is unwilling to give up the presumption that it can continue. As stated in a 2020 *Wired* article on the potential end of the law: "Nonetheless, Intel—one of those three chipmakers—isn't expecting a

funeral for Moore's Law anytime soon. Jim Keller, who took over as Intel's head of silicon engineering in 2018, is the man with the job of keeping it alive . . . He points out that there are probably more than a hundred variables involved in keeping Moore's Law going, each of which provides different benefits and faces its own limits. It means there are many ways to keep doubling the number of devices on a chip—innovations such as 3D architectures and new transistor designs. These days Keller sounds optimistic. He says he has been hearing about the end of Moore's Law for his entire career . . . He says Intel is on pace for the next 10 years, and he will happily do the math for you: 65 billion (number of transistors) times 32 (if chip density doubles every two years) is 2 trillion transistors." See Bruce Sterling, "Preparing for the end of Moore's Law," *Wired*, March 18, 2020, https://www.wired.com/beyond-the-beyond/2020/03/preparing-end-moores-law.

[28] Jessica Matthews, "Early stage startups doubled their funding in 2021," *Fortune*, January 6, 2022, https://fortune.com/2022/01/06/early-stage-startups-funding-doubled-2021. See also "The basics of microchips," Advanced Semiconductor Materials Lithography, accessed September 20, 2022, https://www.asml.com/en/technology/all-about-microchips/microchip-basics.

[29] Clayton Christiansen, *The Innovator's Dilemma* (New York: Harper Business, 1997).

[30] Andy Grove, *Only the Paranoid Survive* (New York: Currency, 1999).

[31] Moving compute power from the desktop to the cloud is explicitly about fast change, including fast mitigation of the flux of vulnerabilities that comes with ever shorter product cycles. The notion of "moving target defense" dictates that if change is too fast to secure, make it faster—the homeopathy of cybersecurity disease, as it were.

[32] "Now, here, you see, it takes all the running you can do, to keep in the same place. If you want to get somewhere else, you must run at least twice as fast as that!" — The Red Queen, quoted in Lewis Carroll, *Through the Looking Glass*.

[33] The rate of change itself being a subcomponent of that environment's overall level of complexity.

[34] Here there is an unavoidable tradeoff between local and systemic risk. See Robert M. May, Simon A. Levin and George Sugihara, "Ecology for bankers," *Nature*, February 20, 2008, https://www.nature.com/articles/451893a.

[35] "Are Optimality and Efficiency the Enemies of Robustness and Resilience?" (the answer is yes, they are). See "Optimality vs. Fragility: Are Optimality and Efficiency the Enemies of Robustness and Resilience?"
(Sante Fe Institute Events Wiki, Annual Risk Meeting, Morgan Stanley World Headquarters, Broadway, New York, October 2, 2014), https://wiki.santafe.edu/index.php?title=Optimality_vs._Fragility:_Are_Optimality_and_Efficiency_the_Enemies_of_Robustness_and_Resilience%3F&oldid=55754.

[36] To those who believe that these issues are either a result of, or best addressed by focusing on, individual actors in a laissez-faire economy, we have only one point to make: that is precisely how things stand now. The dangers we describe throughout this paper are a result of externalities that are not

priced, so to speak, into the transaction costs associated with any given use of software. In other words, caveat emptor has only taken us so far—and look at where it got us. A collective response is needed.

[37] Our hope is that in addressing the latter we may also aid in efforts to fix the former, and that digital risk mitigation might therefore serve as a roadmap for efforts to minimize societal risks writ large.

[38] Repairable? Only if you equate frequent upgrades and attendant feature creep to repair.

[39] See "A Crisis in Third-Party Remote Access Security," SecureLink, Imprivata Company, April 2021, https://www.securelink.com/research-reports/a-crisis-in-third-party-remote-access-security/.

[40] Attacks on supply chains illustrate this point well.

[41] As Nassim Nicholas Taleb describes the complexity cost of interdependence: "[We are] undergoing a switch between [continuous low grade volatility] to . . . the process moving by jumps, with less and less variations outside of jumps." Nassim Nicholas Taleb, "The 'Long Peace' is a Statistical Illusion," https://web.archive.org/web/20121117225617/www.fooledbyrandomness.com/longpeace.pdf.

[42] Some might argue that we have drifted into the politics of inequality, but we disagree. Whether you lean rightward or leftward, there exists shared responsibility for risks and benefits of resources, along with humility about our ability to steer things for the better.

[43] One of us has made a version of this same argument before, with a focus on information security in particular. See James C. Trainor and Andrew Burt, "Our Government's Approach to Cybersecurity Is a Costly Mess. Here's What Would Fix the Problem," Time, January 2, 2020, https://time.com/5757811/cybersecurity-attacks-agency/.

[44] It is worth noting that such recommendations are not entirely new. Indeed, there exist a plethora of past efforts, such as the Trusted Computer System Evaluation Criteria or "Orange Book," published in the mid-1980s by the Department of Defense, intended to increase the quality of software writ large. Decades before that, computer scientists even created a program known as ALGOL, which was intended to reduce errors in programming by becoming the universally used programming language but was never widely adopted. The efforts directed at software assurance over the years, in short, are comprehensive, and we are not advocating that a single such assurance mechanism be used. Instead, we simply advocate that an entity capable of enforcing reasonable standards be empowered to actually enforce such standards on the purveyors of code (the operative word here being reasonable, meaning flexible and adaptable in the legal sense). The hard work of determining what, exactly, constitutes such standards we leave to such an entity itself.

[45] We leave the notion of what, exactly, constitutes a violation to such an entity. Our main point is that there should *be* violations, such that clear standards can be set across the software industry. As but one example, unpatched vulnerabilities in code, beyond a certain period in time after the vulnerability was publicized, might constitute one such violation.

[46] Such as the duty to retain in working order the build environment along with the source code for every fielded version of its product or service.

[47] Here we make an awkward admission: while the digital and the physical worlds are indeed merging, as we argue above, there are core differences between the two. One is the fact that the digital world is, at root, artificial and thereby created by humans, while the other is composed of many elements that are either non- or less-artificial. This should, in our view, not diminish the fact that the two are combining, but neither do we wish to ignore these key differences.

[48] The same does not hold true for the physical world. Who, for example, is responsible for protecting the ocean, the air we breathe, or even the planet? Environmentalism in its physical form presents notoriously vexing issues. For this reason, environmental risks cannot be evenly—or at least easily—distributed so long as responsibilities and ownership over the physical world and its natural resources are not.

[49] Dan Geer, Eric Jardine, and Eireann Leverett, "On market concentration and cybersecurity risk," *Journal of Cyber Policy* 5 (February 2020), https://www.tandfonline.com/doi/full/10.1080/23738871.2020.1728355.

[50] Sandy Clark et al., "Familiarity breeds contempt: the honeymoon effect and the role of legacy code in zero-day vulnerabilities," in *ACSAC 2010: Proceedings of the 26th Annual Computer Security Applications Conference* (December 2010), https://dl.acm.org/doi/10.1145/1920261.1920299.

[51] Dan Geer and George P. Sieniawski, "Who Will Pay the Piper for Open Source Software Maintenance? Can We Increase Reliability as We Increase Reliance?," *USENIX ;login:*, June 2020, https://www.usenix.org/system/files/login/articles/login_summer20_11_geer.pdf.

[52] Another less draconian possibility is simply the public disclosure of environmental impacts for a given code base or device, a tactic that is being used by more traditional environmentalists to raise awareness (and impose reputational costs) on carbon emitters. See "Task Force on Climate-related Financial Disclosures," among many examples, available at https://www.fsb-tcfd.org.

[53] A metaphor we borrow from Andy Ozment and Stuart E. Schechter, "Milk or Wine: Does Software Security Improve with Age?," *USENIX* (August 2006), https://www.usenix.org/conference/15th-usenix-security-symposium/milk-or-wine-does-software-security-improve-age.

[54] Dan Geer, "On Abandonment," *IEEE Security & Privacy* (August 2013), http://geer.tinho.net/ieee/ieee.sp.geer.1307.pdf.

[55] Any device with an IP address becomes a part of the networked digital environment, and therefore it can be used by malicious actors for ill. See Andrew Burt and Daniel E. Geer, Jr., "Flat Light: Data Protection for the Disoriented, From Policy to Practice," Aegis Series Paper No. 1816, Hoover Institution Essay, November 2018, https://www.hoover.org/sites/default/files/research/docs/burtgeer_flatlight_revisednov20_webreadypdf_final.pdf.

[56] This is particularly true as it relates to old devices (think: cell phones, tablets, gaming devices, etc.) with networking connectivity, which must be forcibly retired or able to be patched if we are to preserve some semblance of security in the digital world.

[57] It is the arrival of new threats that marks a code body as needing upgrade, not the code body's time in service.

[58] Effective challenge is defined as "critical analysis by objective, informed parties that can identify model limitations and produce appropriate changes. Effective challenge depends on a combination of incentives, competence, and influence." See Board of Governors of the Federal Reserve System, "Supervisory Guidance on Model Risk Management," SR Letter 11-7 Attachment (April 4, 2011), https://www.federalreserve.gov/supervisionreg/srletters/sr1107a1.pdf.

[59] Andrew Smith, "Using Artificial Intelligence and Algorithms," Bureau of Consumer Protection, Federal Trade Commission, April 8, 2020, https://www.ftc.gov/business-guidance/blog/2020/04/using-artificial-intelligence-algorithms.

[60] See, e.g., Paul Rosenzweig, "Bad code is no excuse, someone is liable," *Financial Review*, August 4, 2017, https://www.afr.com/technology/bad-code-is-no-excuse-someone-is-liable-20170803-gxoxcl.

[61] "Disclaimers and Release," WhatsApp, accessed July 12, 2022, https://www.whatsapp.com/legal/terms-of-service#terms-of-service-limitation-of-liability (note that formatting has been changed for the sake of readability).

[62] There are, of course, some limitations to software developers' ability to avoid accountability, even under existing legal structures. Data breach liability is one such area. Responsibility around hate speech and other harmful content, especially in the European Union, is another.

[63] Including industries that are, or were once considered to be, as innovative as the software industry itself. Innovation, in other words, does not require a blank contractual slate.

[64] See Nora Freeman Engstrom, "When Cars Crash: The Automobile's Tort Law Legacy," *Wake Forest Law Review* 53 (2017), https://law.stanford.edu/wp-content/uploads/2018/06/When_Cars_Crash_-_Tort_Law_Legacy_-_As_Published.pdf.

[65] To make the implications of our arguments more explicit: the introduction of new forms of liability, as detailed in Section 2, combined with the revocation of software vendors' ability to disclaim nearly all forms of liability, will have myriad impacts on such companies' behavior and the use of their products. Once such impact relates to insurance and creation of market incentives to reduce digital risks. Companies are likely to seek to reduce new risks through insurance, whose policies will, in turn, require risk reduction practices to minimize the actual likelihood of downside risk in exchange for coverage, in turn contributing to safer products.