

Cyberensuring Security

Justin (Gus) Hurwitz

Abstract

Cybersecurity is one of the most pressing and legally difficult issues facing this country today. It touches every aspect of modern political and social life, the economy, and national security. From the OPM and IRS breaches, to the Sony hack, to attacks on hospitals and health insurers, to attacks on domestic and international infrastructure, to domestic and international surveillance, cybersecurity concerns are omnipresent. For technical, legal, and practical, reasons, they also have proven extremely difficult to address.

This article draws from the economic literatures on strict liability and insurance to argue that cyber incidents generally, and data breaches in specific, should be treated as strict liability offenses. But that is only the starting point of this article's argument. The economic literature on strict liability recognizes that it is, in fact, a form of insurance – potential tortfeasors subject to strict liability effectively are required to insure others against harms caused by their conduct. This article's core argument is that pervasive cyber-incident insurance is the best approach to addressing the full range of cybersecurity concerns.

The characteristics of the model proposed in this article compare favorably to the current status quo – one in which users are largely helpless, firms are largely unknowledgeable, software is generally insecure, federal agencies are generally impotent to bring about meaningful change, and attackers are largely judgement-proof. As an initial matter, it would offer consumers redress when cyber-incidents occur. But more important, it would facilitate education about and monitoring of cybersecurity practices; it would facilitate the collection, analysis, and use, of aggregate information about the causes and costs of these incidents; and it would put that information the hands of parties in a position to improve the existing ecosystem.

Table of Contents

- I. The Cybersecurity Challenge
 - I.A. The *ex ante* Cybersecurity Challenge
 - I.B. The *ex post* Cybersecurity Challenge
 - I.C. The Multiplicity of Actors
 - I.D. The Multiplicity of (conflicting) Incentives
- II. The Current (and Ineffective) Legal Approaches to Cybersecurity
 - II.A. Law and Technology as Complementary Approaches to Cybersecurity
 - II.B. Private Law Approaches to Cybersecurity
 - II.C. Public Law Approaches to Cybersecurity
- III. Strict Liability for Cyber-Incidents: The Sword
 - III.A. Defining Strict Liability
 - III.B. Cybersecurity is a Classic Case for Strict Liability
 - III.C. Limitations: Statutory Damages, and Other Practicalities, and Best Laid Plans
- IV. Data-Breach and Cyber-Incident Insurance: The Shield
 - IV.A. Strict Liability as Insurance
 - IV.B. Insurance as Regulation
 - IV.C. Jumpstarting the Market

Cyberensuring Security

Justin (Gus) Hurwitz

[Discussion draft; Footnotes omitted]

Introduction

Securing Internet-connected computers and information stored on those computers is legitimately hard. On one side of the equation, any Internet-connected computer affords any number of potential attackers any number of paths to compromise that computer for any number of reasons. On the other side of the equation, it is unclear who can or should be responsible for taking action to prevent, mitigate, or respond to an attack. Users are rarely sophisticated enough to understand, let alone meaningfully take action to address, computer security – and they are wholly unable to protect their information once it is in the hands of a bailee third party; it is remarkably difficult to design “secure” computer software (indeed, it is mathematically provable that any program beyond a trivial level of complexity cannot be provably secure); providers of software and communications infrastructure are generally immunized by contract or law, and in any event are generally too remote from specific incidents to be held liable for harm; private law mechanisms have proved wholly inadequate to address security incidents – even if a victim can identify a tortfeasor from whom recover is possible, proving causation and harm is difficult; and public law institutions simply lack the resources to meaningfully stem the rising tide of security incidents.

This article draws from the economic literatures on strict liability and insurance to argue that cyber incidents generally, and data breaches in specific, should be treated as strict liability offenses. This stands in contrast to the approach taken to date in the United States, which can be characterized as a mish-mash of sector-specific regulations backed up by a largely ineffective effort by the Federal Trade Commission to develop negligence-like general data security norms. The challenges of ensuring cybersecurity suggest that a strict-liability approach is more appropriate to the nature of both the risk and the potential harms.

But that is only the starting point of this article’s argument. The economic literature on strict liability recognizes that it is, in fact, a form of insurance – potential tortfeasors subject to strict liability effectively are required provide insurance for those who may be harmed. Indeed, it is often the case that entities subject to strict liability purchase third-party insurance on behalf of those who may bring claims against them. Following in this vein, this article’s core argument is that pervasive cyber-incident insurance is the best approach to the full range of cybersecurity concerns.

While surely not a silver bullet solution – no such solution exists – this approach has many advantages. First, drawing on the theory underlying strict liability, it places the burdens of avoiding and redressing harm on the relatively sophisticated and least-cost party in the consumer-firm relationship. Second, it creates substantial incentives for the rapid development of a robust market for security insurance products. To date, this market has been slow to develop – and, to the extent that it has developed, it has focused on exposure by large firms to coarsely-defined incidents. The establishment of a robust insurance market will have myriad beneficial secondary effects. Through the claims process, insurers are natural, neutral, data aggregators – data breach insurance will have much the same effect of information sharing legislation recently considered by Congress. Through the underwriting and renewal process, insurers can educate users and firms about good security practices, and audit firms for compliance with those practices. And insurers also overcome collective

action problems that have historically advantaged some participants in the information economy – especially software developers and online intermediaries, both of whom are substantially shielded from liability today, and the actual malfeasing attackers who in the civil context rarely face sophisticated adversaries with the ability or incentive to pursue them.

To facilitate the turn to pervasive cyber-incident insurance, this article proposes relatively simple statutory approach. This approach is based upon a two-stage liability regime. The first stage is the creation of a new federal cause of action imposing strict liability upon entities that experience cyber-incidents resulting in harm to third parties. The Federal Trade Commission (or another agency) would be empowered to define both specific types of incidents and harms which would trigger this liability, and also specific statutory damages for these harm (based upon reasonable empirical evidence) – though it would not necessarily have enforcement authority of its own. The second stage of this liability regime provides both a carrot and a stick in favor of cyber-incident insurance: firms that *do not* have effective cyber-incident insurance policies at the time of an incident would be subject to a statutory damages multiplier; firms that *do* have effective cyber-incident insurance, on the other hand, would only be subject to actual, provable, damages. In this latter case, the FTC's schedule of statutory damages would be persuasive evidence to be considered by a court in determining damages, but insurers would have the opportunity to contest that schedule based upon their own empirical actuarial data.

The characteristics of this model compare favorably to the current status quo – one in which users are largely helpless, firms are largely unknowledgeable, software is generally insecure, federal agencies are generally impotent to bring about meaningful change, and attackers are largely judgement-proof. As an initial matter, it would offer consumers redress when cyber-incidents occur. But more important, it would facilitate education about and monitoring of cybersecurity practices; it would facilitate the collection, analysis, and use, of aggregate information about the causes and costs of these incidents; and it would put that information the hands of parties in a position to improve the existing ecosystem.

This article proceeds in four parts, outlined below.

The article begins in Part I by outlining the basic challenges facing the cybersecurity landscape. This discussion starts with the technical difficulty of designing and implementing a secure system. It is, simply, hard to do right. Once a system is implemented, it is then extremely difficult to monitor that system in a way to effectively detect intrusion; once an intrusion is detected it is difficult to mitigate its ongoing effects in the short run; and it is even harder to effectively respond to them in the long-run. And even if we *could* overcome these largely technical difficulties, it is unlikely that we actually *would* do so, due to the large number of actors involved in the security ecosystem and their myriad, and oftentimes conflicting incentives.

The discussion turns to the legal challenges of cybersecurity. Part II looks at the difficulties that have faced – really, plagued – both private law and public law efforts to address these concerns. Courts have effectively made it impossible to bring suit over cybersecurity-related harms. From declining to treat software as a product, to allowing its sale to be governed by indemnifying contracts and licenses, to refusing to find probable causation or cognizable harm – at every step of a possible lawsuit the courts have made it very difficult for a legal challenge to survive. On the public law side, various regulatory efforts to improve cybersecurity have largely failed. While they effectively punish a small subset of firms that experience cyber-incidents, they have done very little to improve the overall quality of the cybersecurity ecosystem. Moreover, with the exception of the FTC, most regulatory efforts have focused narrowly on specific industries – and the FTC's efforts

are largely inadequate to regulate the vast majority of entities subject to potential cybersecurity-related harms: individuals and small and mid-size business.

Having explained why the cybersecurity landscape is so difficult, both to navigate and to improve, Parts III and IV turn to a solution. Part III argues that a transition to strict liability for cybersecurity-related harms would remedy the majority of the concerns raised in Parts I and II, allowing standard private law institutions to function and bring about dramatic improvements to the cybersecurity ecosystem. In addition, Part III offers some suggestions for how to implement such a transition, addressing specifically the need for statutory damages to accompany strict liability and suggesting the Federal Trade Commission be tasked with establishing a schedule of damages.

Transitioning to strict liability would address many current cybersecurity difficulties. But there is another, more powerful, benefit to strict liability: strict liability is an insurance regime. Part IV argues that broad adoption of strict liability would foster the development of a vibrant market for data-breach and other cyber-incident insurance policies. This offers an approach to addressing cybersecurity concerns that is substantially better than any liability-based model. This is because cybersecurity-related risks are based in systemic problems that affect users throughout the cybersecurity ecosystem – they rarely result from individual-scale problems. Such risks are better managed on a pooled, insurance-scale basis. And, importantly, insurance-based systems have powerful regulatory characteristics that can both improve the behavior of individuals operating in the current cybersecurity ecosystem and improve the overall state of that ecosystem moving forward.

This article concludes with a brief summary of its arguments, structured to highlight key elements as a coherent policy proposal.

I. The Cybersecurity Challenge

Computer security, especially for Internet-connected devices, presents legitimately hard problems. This section offers a brief overview of the technical, practical, and some legal, difficulties of protecting data in the online environment.

I.A. The *ex ante* Cybersecurity Challenge

The basic task of cybersecurity is seemingly simple: to allow that those, and only those, authorized to access data or computers systems are allowed to do so. The difficulties of accomplishing this easily-stated task, however, are myriad. We can begin with just scoping the elements of the task: we need a way to specify who does and does not have access to a secured system, a way to identify those users, and a way to specify and control the level of access they may have. This alone decomposes into requirements to specify and authenticate individual users, specify the various resources they may be able to access, and specify various levels of access each user is allowed to each resource. The number of combinations of users, resources, and permitted access levels permutes quickly, imposing substantial costs (mostly in the form of managing a complex system) both on those managing and those using the secured system.

Indeed, this complexity is one of the fundamental trade-offs in the world of security: allow for finer-grained control, which increases complexity (imposing higher burdens on those subject to the security model and also increasing the likelihood the mistakes in implementing that model will be made, leading to security faults), or reduce the complexity of the system and need to either allow some users greater access to secured resources than is necessary or deny some trustworthy users access to resource of which they would otherwise be able to make beneficial uses. By analogy, one could imagine putting locked doors, each with a unique key, on every room in an office building, and providing employees with individual keys to each room to which they are allowed access – but no office would actually operate that way, because it imposes such great costs on both the office management and individual employees. And this also demonstrates another of the fundamental trade-offs in designing secure systems: if the cost of complying with security protocols is too great, users may find ways to bypass those protocols. This, too, is most easily seen by example: in the physical world, one may prop doors open instead of continually unlocking them; in the world of computers, users may leave passwords written on post-it notes, or use the same password for all of their online activity. [[[Discuss hygiene, cite reports that users are most common vector; discuss that users value usability, don't see security, so market provides usability over security]]]

The challenge created by complexity is much more problematic in the case of computer security than in that of the physical world. This is because every aspect of computer security needs to be implemented in computer code. There are two basic reasons that this is difficult. First, it requires every aspect of the security protocols to be *completely* specified *ex ante*, and, second, these security protocols must be specified (in computer code) *accurately*. Again taking the physical world as a counter-example: one need not “program” a door to allow the fire department in in the case of an emergency (either humans will intervene and, smelling smoke, will allow entrance, or the door will be forcibly circumvented); one need not “program” employees to comply with a court order or warrant; and individual actors can accommodate many otherwise incompletely-specified actions on an ad-hoc basis (for instance, an employee could make photocopies for a contractor who does not have a copier code). In the computer context, each of these actions would need to be specified *ex ante* – otherwise the system may need to be taken offline and reprogrammed, or otherwise circumvented on a case-by-case (and likely complex) basis. Of course, we could imagine implementing computer-based security protocols in a way that allowed for greater human discretion (e.g., a bank teller could review every online-transaction a user makes, or a system administrator could confirm a users’ credentials each time she logs into a system). But doing so would defeat one of the basic purposes of computer-based interactions: removing humans from routinized transactions so that those transactions can proceed at computer-scale, not human-scale, speeds.

The problem of *accurate* implementation is even more substantial than that of complete implementation. Indeed, one of the foundational theorems in computer science – the so-called “halting problem” – effectively states that it is effectively impossible to prove that any computer code beyond a trivial level of complexity operates as intended (that is, that it contains no bugs, such as those that could render a security protocol ineffective). We need not delve into the mathematical proofs that it is nearly impossible to prove that a given piece of computer code is defect-free. Rather we can point to some of the canonical examples of basic implementations mistakes in security related software. Examples include the Heartbleed bug, bugs in the Apache TLS implementation, attacks on encryption Certificate Authorities, the Shellshock bug, critical security vulnerabilities in

the Linux operating system and related programs, encryption flaws in Apple iMessage, and LastPass. Each of these is an example of code that has been developed and scrutinized, often for years, by sophisticated, security-conscious, programmers that nonetheless contained critical flaws in how they were implemented. In other words, at a technical level security is hard – very hard – to do correctly.

I.B. The *ex post* Cybersecurity Challenge

The issues discussed above relate to the challenges of designing and implementing a system that is secure – that is, a system that prevents unauthorized activity. This is, however, only part of the cybersecurity challenge. Because no system is completely secure, any sound security design needs to anticipate and respond to security breaches. Incident response presents its own slate of problems, including technical challenges similar to those that make designing and implementing secure systems difficult, physical-world problems relating to coordinating human resources to respond to incidents, and legal challenges.

As an initial matter, systems need to be designed to allow for the detection of security breaches. Here, as above, the task of programming computers for this task is far more difficult than analogous physical-world challenges. To start, secured systems need to be instruments with monitoring capabilities that can observe and record how they are used. This is an onerous, and at times intractable, task. Adding such capabilities can substantially reduce system performance, such that any monitoring instrumentation needs to be deployed sparsely. We also face the same challenge of implementing it correctly. Attackers therefore already have two attack vectors: attack resources that are either unmonitored or ineffectively monitored. What is more, it is frequently the case that an attacker who breaches a secured system simultaneously (or as a result of the breach) obtains access to the system’s monitoring capabilities. Generally, avoiding this conundrum requires implementing additional, separate, monitoring and logging systems (that is, computers) – but this has the unfortunate consequence of increasing overall system complexity even further, which can actually make it easier for breaches to occur!

One function of well-implemented monitoring tools is to detect security incidents in real time. But monitoring also serves the important function of recording system activity for later use and analysis. The simplest aspect of this to understand is allowing for the reconstruction of incidents. Reconstruction serves at least three important purposes: to figure out how an attack occurred so that future attacks can be prevented, to understand to effects of the attack (e.g., to see what data was compromised), and to serve as evidence in identifying and taking action against those responsible for the attack.

But there is another, as important, purpose behind monitoring: establishing a baseline of “normal” system operation. Unfortunately, data is rarely analyzed for these purposes. This is one of the reasons that the second part of the monitoring equation is rarely satisfactorily met: the ultimate purpose of *monitoring* a system is to *detect* abnormal behavior.

The seemingly key function of detecting attacks is often the most challenging to accomplish. In most computer security breaches investigated by security consulting firms, the attackers breach a system several months before their breach is detected. In this time, they may be engaged in

malicious activity (such as exfiltrating information, or manipulating internal information to harm an attack target), or they may be using their initial breach as a beachhead to further penetrate a target's systems.

Once a breach has been detected, incident response becomes the order of the day. The first step in incident response is to *mitigate* any ongoing effects of the breach. For instance, compromised systems should be disconnected from any networks, sensitive accounts should be locked down, and appropriate resources should be engaged (e.g., law enforcement or security professionals.). Here too proper response can be both technical and difficult. For instance, one of the most important things to *not* do upon discovering a compromised system is to turn the system off – even though this is the intuitive response. Turning the system off will delete potentially important information stored in the computer's memory and terminate any active programs – information that can be used to figure out the source and scope of an attack. And turning a system back on can overwrite similarly important information.

Once the effects of the incident have been mitigated, the compromised party can turn to responding to the attack. This may include any number of efforts. For instance, a compromised system should be repaired, and the source or cause of the compromise fixed to prevent future incidents. Parties harmed by an attack should generally be notified, both as a matter of best practices and often as a function of relevant law. Compromised data or systems may need to be replaced or repaired. The victims of an attack may want to work with law enforcement, insurers, or vendors to identify or take action against the attackers. And the victims of an attack may want or need to take legal action of their own (e.g., to bring a civil suit against their attackers, if possible, or to defend themselves against suits brought by the government or as a class action). Discussion of the viability, practicalities, and limitations of such legal action are the subject of Part II.

I.C. The Multiplicity of Actors

The technical difficulties of designing secure computer systems are dramatically compounded by the sheer number of actors in the security ecosystem. It is useful to identify these actors here, before discussing how their (often conflicting) incentives further complicate computer security.

On one far side of the web, we have “users” – those who actually use a (possibly) secure system. Even this basic unit of the ecosystem is more complicated than one would expect. “Users” can refer to the consumer end-users of a piece of software, such as Microsoft Windows. In a firm, users may refer to the employees of the firm who use software purchased or designed by the firm. The firm itself may be said to be the user of software purchased by use by its employees, or even of software that it designed (or bespoke software designed by a contractor). And, of course, the firms' customers are “users” of services offered by the firm, which may or may not rely upon systems designed or implemented by third parties.

Those third parties may be firms such as Microsoft, Apple, Google, or ExamSoft. They could also be vendors that sell turnkey solutions, designed by themselves or by others. They may be contractors, who design bespoke systems. Or they may be “integrators,” who integrate various

platforms designed by third parties with a firm's own systems. Connecting all of these systems are various Internet-based entities. This includes the ISPs that connect firms to the Internet, or a firm's customers to the firm's servers and services. It also includes cloud-based services, which often host information.

This multiplicity of actors makes establishing cybersecurity responsibility difficult. Each of the actors has some legitimate argument that at least some of the others bears responsibility almost any cyber-incident. Software was poorly designed or integrated; systems were improperly implemented or managed by firms or their contractors; ISPs should have detected harmful activity by malicious actors and informed their targets or cut off access; firms failed to train or monitor their employees, or employees failed to comply with established procedure; customers chose to work with unknown firms that had unknown or poor security practices, or firms failed to have satisfactory security practices. The response to any security incident will invariably be to assign blame to any number of other parties.

It is notable that responsibility for security breaches is rarely meaningfully attributed to the parties that are actually responsible: the attackers behind the cyber-incident. This is a nod to the practical reality that it is often impossible to identify the attackers, that it would typically be almost impossible to bring suit against the attackers if it were possible to identify them, and that even then one would be unlikely to recover meaningful damages from them. Amazingly enough (and as discussed in Part II), even if you could find the attacker, given the challenges discussed throughout this Part, it would be very difficult to establish the elements required to be awarded damages against them – given the multiplicity of actors and difficulties of designing secure systems, it can be difficult, if not impossible, to establish causation and (especially) harm.

On a final note, it is useful to discuss briefly the multiplicity of harms that may result from a cyber-incident – or, stated alternatively, the multiplicity of motivations that attackers may have. Starting with the most apparent motivations that attackers may have: they may seek to obtain information through hacking. This could be information about a firm's customers (e.g., passwords, personal information, correspondence, credit card information), or about the firm itself (as in the case of espionage). They may also intend to damage a firm, for instance by altering or deleting sensitive information, or damaging physical systems controlled by compromised computers. Attackers may use “ransomware” to demand money from a target. The compromised systems may in fact not even be an intended target: they could be a platform that attackers use in attacking other, third party, systems. Or attackers may have political or social purposes: they may intend to embarrass a firm or individuals, to cause reputational damage, or to advance a political agenda. This range of motivations further demonstrates the challenges discussed so far. For instance, if one expects attackers to target sensitive customer information, it may be possible to address this risk by minimizing the amount of customer information that is stored and encrypting what information must be kept. It is more difficult, however, to prevent control systems from being used to damage the systems that they control – to do so undermines the purpose of having computerized control systems. And this also demonstrates what will be an important challenge discussed in detail below: establishing harm for the purposes of liability. How should a court measure the harm caused by an attack that shuts down a firm for a few hours, or results in the disclosure of (truthful) information about a firm's customers, or that is the basis forcing the firm to adopt a new policy as part of a

political agenda? Courts are generally reluctant to award damages for harms such as these – they are simply too speculative and difficult to measure.

I.D. The Multiplicity of (conflicting) Incentives

Each of the myriad actors in the cybersecurity ecosystem faces their own incentives in deciding how – or whether – to respond to security concerns. While each would likely benefit from an improved cybersecurity ecosystem, none has strong incentives to invest substantially in such benefits. And many, in fact, have incentives to adopt but security practices.

Perhaps the most important set of incentives echoes the fundamental trade-offs between security and usability described in Part I.A. Both users and those designing software and other computer systems are generally willing to forego security for greater usability and performance. [[[Cf. addition of “users want usability” discussion in I.A.]]] This results in large part from the difficulty of holding designers liable for defective software – the threat of legal liability would of course be a powerful incentive for firms to improve their products’ security. This is further exacerbated by firms’ ability to attribute fault for security incidents to others in the security ecosystem – including attributing fault to users themselves. This reduces firms’ ability (or need) to compete along a security dimension – especially when consumers are often more responsive to the usability and short-term cost dimensions. And there is reason to argue that users do, in fact, bear some responsibility for the poor state of the cybersecurity ecosystem: despite professed fears about the collection and use of sensitive data, and widespread concern about cybersecurity, consumers very readily engage in conduct online that exposes them to risks. This is surely, in some part, a reflection of putatively irrational decisionmaking by consumers. It is also, to some extent, a form of rational ignorance: consumers are not security experts, they do not have the time or knowledge necessary to evaluate most firms’ security practices, and they reasonably believe that the law will protect them should they be harmed by malfeasant firms – so it is eminently reasonable for consumers to engage in what appears to ordinary users to be ordinary online activities.

Similarly, firms that make use of third party security systems in their broader business – that is, the vast majority of firms – face poor security incentives. Most security incidents target firms’ customers’ information, such that the firms themselves are unlikely to experience any loss from an attack – unless, that is, the fact of the attack becomes public, in which case a firm may face substantial reputational harms and may also bear some direct costs from responding to the attack. In other words, the incentive for most firms is to invest in very basic security – only enough to secure their systems from casual attackers – and otherwise pay no attention to security. These incentives were arguably even worse before the recent, and relatively widespread, adoption of state data breach notification laws – laws that require firms to notify affected consumers of data breaches that may affect their (the consumers’) data. Before the adoption of such laws, firms often faced no incentive to disclose, or even to respond to, a data breach, and faced incentives to keep the fact of the breach secret. But even following adoption of these laws, firms still do not face substantial incentives to adopt strong security practices. This is in part because there is still little likelihood that a firm will be held liable for damages resulting from a data breach. More tragically, this is also largely because

consumers have become inured to data breaches, such that the reputational harm to a firm of a data breach is much less today than it was even two or three years ago.

Perhaps the worst incentives are faced by the cybersecurity industry itself. Estimates vary, but the size of the cybersecurity “market” – comprising firms that specialize in various aspects of cybersecurity, from systems design, to consulting, to incident response and litigation – is currently pegged at somewhere around \$75-100 billion. This amount is expected to grow to \$150-170 billion by 2020, a growth rate significantly exceeding that of other parts of the economy. In other words, the status quo is working well for this industry. Its participants have little reason to improve the state of the cybersecurity ecosystem.

II. The Current (and Ineffective) Legal Approaches to Cybersecurity

Part I looked at why cybersecurity is difficult as a technical and practical matter. This Part looks at the difficulties of using the law to address cybersecurity concerns. It starts by considering the relationship between legal and technical institutions. It then considers the challenges that private law institutions have faced in responding to cybersecurity incidents, followed by consideration of the efficacy to date of public law institutions.

II.A. Law and Technology as Complementary Approaches to Cybersecurity

As discussed in Part I, there are many reasons that cybersecurity is a legitimately hard problem. It is technologically difficult to specify what is required of a secure system, it is extremely difficult to accurately implement the system once specified, and it is effectively impossible to verify that such a system is implemented correctly. Moreover, the costs of security – in terms of design, implementation, performance, and user experience – are substantial enough that they are often not justified by their benefits. This is particularly problematic when we consider the private incentives faced by almost every actor in the security ecosystem: almost no actor has strong incentives that are in line with best security practices, and almost every actor has strong incentives that run contrary to best security practices.

In other words, security is hard.

None of these problems, however, is new. Many systems and institutions are difficult to design or implement properly. Indeed, it is a basic fact of life that mistakes and accidents happen, and that people are often harmed by those mistakes. Moreover, it is very often the case that individuals’ private incentives do not line align with socially-optimal conduct. Cybersecurity presents extreme cases of all of these problems.

Society manages to continue moving along despite these challenges. This is largely because the law operates as a backstop that mitigates the harms that may result from them. The law steps in where things go wrong. In general it does so through two mechanisms. First, it compensates parties that are harmed by bad actors or bad actions. In other words, it assures users of a system that if they are acting reasonably and are harmed by another actor who is acting unreasonably, that they can be compensated for that harm. And, second, it makes clear that parties who cause harm to occur are

liable for that harm. This, in turn creates incentives for those who create systems used by others to do so carefully – to design their systems so that they will not cause undue harm, because they will be responsible for compensating others for those harms.

Law and technology are complementary approaches to the design of well-designed systems. We want systems to be well designed ex ante so as to limit harms. The availability of ex post remedies ensures that those designing systems will be held to account for their design decisions. At the same time, the law recognizes that risk is inevitable – that reasonable mistakes may happen. So the law generally works to assign liability for harms in ways that maximize the social value of activity, mitigating concerns that individual actors will be motivated solely by their private incentives at the expense of imposing costs on society.

But this synergy between law and technology assumes effectively designed and implemented legal rules. As suggested in Part I, the legal rules relating to cybersecurity have proven wholly ineffective. This has had the unfortunate consequence of exacerbating cybersecurity problems – as described in Part I, many actors in the cybersecurity ecosystem not only lack incentives to act well, but have incentives to act badly. These incentive mismatches result largely from the lack of effective legal rules.

The rest of this Part discussed the failings of current legal approaches to addressing cybersecurity concerns. This will provide a foundation for the discussion in Parts III and IV, presenting an alternative approach to these concerns.

II.B. Private Law Approaches to Cybersecurity

“Private law” refers broadly to legal causes of action that individuals are able to bring against one another. For instance, suits for trespass, breach of contract, or negligence are traditional private law causes of action. So too would be a civil cause of action created by statute that can be initiated by individuals, or a class action brought by a group of individuals. This is in contrast to “public law” causes of action, which are generally those initiated by the government. These include, for instance, criminal prosecutions, enforcement actions brought by regulatory agencies, rules created by federal agencies with which regulated parties must comply, and various forms of informal regulation exercised by government actors to channel the conduct of private parties.

Private law institution have proven largely ineffective at addressing cybersecurity concerns for much the same reason that cybersecurity is itself difficult. In order to successfully bring a civil lawsuit, one needs to be able to demonstrate various things, such as the identity of the actors that caused a harm; that they, in fact, did cause that harm; that the harm is legally cognizable; and that there is some adequate measure of damages. Each of these elements is difficult in the context of cybersecurity. The multiplicity of actors in the security ecosystem and the complexity of the interactions between them makes it difficult, and sometimes impossible, to attribute fault to any specific actor. And even when fault can be attributed to a single actor, there are likely other confounding factors (or actors) that make it difficult to prove that that actor’s conduct was a proximate cause of the specific harm. For instance, a firm that failed to safeguard its customers’ information may argue that the software it was using was defective, that the vendor hired to install

and maintain its software failed to do so correctly, that an auditing firm it hired to ensure its systems were properly secured failed to detect the relevant faults, that its network providers failed to detect or alert it to suspicious activity, or even that the customers were contributorily negligent in providing their data to an untrustworthy party.

Even if the harmed party can demonstrate that a specific actor's conduct was improper and proximately caused an adverse security incident, courts have struggled with the concept of "harm" online – both in terms of recognizing that the subject of the cyber-incident has in fact experienced harm and in assessing the extent of that harm for purposes of damages. The canonical example here is the disclosure or theft of personal information. In one canonical case, for instance, courts found that an airline's disclosure – in violation of contractual assurances prohibiting such disclosure – of passenger information to the federal government's anti-terrorism efforts didn't represent a cognizable harm to the customers. In that case, the court dismissed the lawsuit because the lack of awardable damages rendered it moot. Similarly, courts have struggled with cases of identity theft or theft of credit cards, especially where credit monitoring services are provided to affected customers or banks refund fraudulent charges. And even where courts are willing to recognize that harms are real, the question often turns back to questions of proximate cause: we live in a world in which information such as credit card numbers is stolen with such frequency that it is difficult for a court to accept that fraudulent charges resulted from any specific theft of a consumer's information – it is simply too possible that the specific harm resulted from some other cyber-incident for the courts to award damages against a possibly-innocent third party without some greater evidence tying the fraudulent use of credit card information to a specific breach. Of course, such evidence is almost certainly impossible to gather.

There is another issue lingering in the background of the discussion so far. The sort of cases discussed above – in which a firm fails to properly protect its customers from adverse cyber-incidents – are governed by tort law, specifically negligence. Other forms of tort claims are similarly problematic (e.g., intentional torts, such as trespass, which as a firm may want to bring directly against attackers, are problematic because it is very difficult to identify the attackers, to attribute a specific attack to them, prove causation, demonstrate no contributory factors that offer the attackers a defense, and demonstrate cognizable, recoverable, damages). But other issues are governed by contract law. Contract law is important in the cybersecurity context for two critical reasons. First, courts have generally upheld the use of contracts – including the sort of dense, boilerplate, consumer-facing contracts that are widely recognized as meaningless to consumers – in the cyber-domain. The contracts very often contain waivers of liability or other forms of indemnification. Unfortunately, liability is typically contractually assigned away from parties that are most likely to be provable liable, or otherwise limits damages. This further compounds the problems of determining liability discussed above. And second, contractual language is often imprecise – a reflection of the complexity inherent in the cybersecurity ecosystem – which in many cases creates further uncertainty rather than clarifying responsibility.

Importantly, private law has a relatively simple mechanism for dealing with many of the difficulties that have been discussed: strict liability. Under some circumstances the law will assign liability to a given party regardless of fault. The canonical area of strict liability in tort law is products liability: the manufacturer of a defective product that causes consumer harm is liable for any harms

caused by that product no matter how negligently the consumer was in its use. Thus, for example, the manufacturer of a table saw would be liable for a injuries caused to a consumer by a failure of the saw, *even if* the consumer were using the saw for improper purposes, while intoxicated, after damaging the saw, and while wearing a blindfold and standing on crutches. Or, as another example, someone who chooses to engage in “ultrahazardous” or otherwise extreme activities – blasting with dynamite, or keeping dangerous animals like tigers as pets – is generally subject to strict liability.

The underlying policy rationales for strict liability are discussed in Part III, which argues that cybersecurity should be a strict liability regime. For the purposes of the present discussion, we need only say that courts have declined to treat services or computer software – the primary components of the cybersecurity ecosystem – as “products.” They therefore have not been treated subject to the rules of strict liability. Rather, they have been subject to the traditional principles of contract and negligence.

II.C. Public Law Approaches to Cybersecurity

There are various public law institutions in the United States that address cybersecurity issues. While some of these efforts effectively address narrow problems that effect parts of the cybersecurity ecosystem, there are no effective public law institutions that address broader problems. In particular, there are no public law institutions that generally ensure parties harmed by adverse cyber-incidents can secure recovery for their losses, that alter the perverse incentives faced by the various actors in the cybersecurity ecosystem, or that generally improve the overall quality of that ecosystem.

In the United States we have no general law of data security. Rather, we have taken a sector-by-sector approach to regulating specify security concerns. There are, for instance, specific laws relating to the security of financial information, health information, information about students, and consumer credit information. Certain industries are also subject to security-related regulation, such as the energy and communications industries.

By and large, regulatory efforts to improve security such as these are inoffensive. Without question they draw additional attention and scrutiny to particularly sensitive areas and provide valuable resource towards the goals both of educating stakeholders about security concerns and of taking action against those who fail to address these concerns. At the same time, we should be aware of the limitations of targeted approaches such as these. In almost every instance, sector specific regulations are “consumer protection” statutes that impose strict controls on what information can be shared or used by those to whom it has been given. Firms generally implement these requirements by substantially limiting how information they hold can be access by employees or shared among their peers or partners. While this has the positive effect of protecting consumers, it has adverse effects of limiting the use of more efficient technologies or making more valuable uses of information. For instance, restrictions on the use and sharing of medical information dramatically hampers medical research – it is literally the case that medical researchers believe that we would have already cured many forms of cancer if not for HIPAA. Restrictions on financial transactions and disclosure of student records encourage firms to use outdated systems, impose burdens on

consumers who need to authorize the disclosure or use of their information, and generally lead industry to make use of stale, but statutorily-clear, business practices instead of innovating new ones.

More problematic, because these rules are generally focused on protecting consumers, they are not focused on improving the overall state of the cybersecurity ecosystem. As such, they don't offer a systematic approach to addressing any of the issues that make cybersecurity difficult. Because these regulations are industry-specific, but the issues that make cybersecurity hard are generalized, none of the regulated industries is in a strong position to effect change to the broader cybersecurity ecosystem. Rather, each industry develops its own, costly, and largely inefficient (if not ineffective) means to protecting consumers.

[DISCUSS CISA]

The Federal Trade Commission is the great exception to the sector-specific approach to cybersecurity in the United States. Since the turn of the century, the Commission has been working to use its general authority to regulate “unfair and deceptive acts and practices” under Section 5 of the FTC Act to establish itself as a general regulator of consumer-facing data security issues. The FTC got into the business of regulating firms' data security practices largely in response to the failure of the private law described above. After courts began dismissing lawsuits because consumers could not establish harm, the FTC stepped in to take action against firms accused of mishandling consumer data, arguing that failure to protect consumer data was an unfair (or, if in violation of a firm's established security or privacy policy, a deceptive) business practice.

The FTC's efforts have been controversial, both lauded and criticized by many. Much of the controversy over the FTC's efforts relate to its use of broad and uncertain legal authority to regulate an large portion of the economy without clear Congressional authority to do so, and in particular its use of adjudication (as opposed to rulemaking procedures) to develop binding legal norms. What this means is that the Commission has not provided the industry with any formally-issued guidance regarding what constitutes “good” or “bad” security practices. Rather, it has offered informal guidance on an occasional, often ad-hoc, basis, which it has sought to formalize by taking legal action against firms that, in the FTC's own estimation, are engaged in bad behavior. Due to the procedures the FTC has used in approaching this issue, the legality (and constitutionality) of this approach has not been addressed by the courts – though one case is currently pending that may lead to such a resolution.

Regardless of the legality of the FTC's efforts to regulate data security practices, there are other reasons these efforts should raise concern. As an initial matter, the FTC approaches security from a consumer protection perspective. As such – and as with the sector-specific approaches – its efforts focus only on the outer border of the cybersecurity challenge. The FTC does not try, not does it have statutory power to try, to address the myriad actors and mixed incentives that make ensuring cybersecurity difficult. It is possible that the FTC's approach will, over time, indirectly influence the incentives of the myriad actors in the broader cybersecurity ecosystem: as firms become increasingly aware that they may face liability for failure to protect consumer data, those firms may demand more secure systems from the rest of the ecosystem. This effect, however, will likely be largely muted in the case of the FTC's enforcement actions. As an initial matter, firms may choose, instead, to adopt clear policies indicating that consumers use their services at their own risk,

or otherwise limiting their liability. Indeed, on the FTC's own terms its efforts are only meant to hold firms to "reasonable" security practices, which should arguably be weighed in light of the current state of the art – these efforts therefore ought not to create any incentives to change the state of the art on their own.

Another important problem with the FTC approach to cybersecurity is that does not meaningfully inform or educate anyone about good security practices. The primary audience for the FTC's data security are a small cadre of data security lawyers and information security professionals who work at relatively sophisticated, mid- to large-sized, firms. This further insulates the effects of the FTC's efforts from the core cybersecurity challenges. First, to the extent that it is educating firms about good cybersecurity practices, the FTC is only communicating to those firms that already understand the challenges of cybersecurity, and that largely have the internal resources to address these challenges on their own. But the vast majority of online activity is undertaken by less-sophisticated actors – consumers, small businesses, and start-ups, who either lack sophisticated understandings of, or the resources to address, cybersecurity challenges. And, importantly, these are the same actors who depend on outside resources – the myriad parties with mixed incentives that permeate the cybersecurity ecosystem – to educate and protect them.

III. Strict Liability for Cyber-Incidents: The Sword

Part II explained that the law, when working well, can create powerful incentives that align individual conduct with socially-optimal goals – but that, in the case of cybersecurity, various factors confound the law's utility. This Part argues that a transition to strict liability for cybersecurity-related harms would remedy the majority of these concerns, thereby allowing standard private law institutions to function and bring about dramatic improvements to the state of the cybersecurity ecosystem. In addition, this Part offers some suggestions for how to implement such a transition. Importantly, this is only the first part of this Article's broader recommendation – in Part IV we will turn to the desirability of a vibrant cybersecurity insurance marketplace and the relationship between strict liability and such a marketplace.

III.A. Defining Strict Liability

The primary private law mechanism that has been used – or attempted to be used, as described above – to address cybersecurity concerns is negligence. Under this model, parties are only liable for harms that they cause to others through their own negligence. In the classic formulation, a party engaging in an action that causes harm to someone else has acted negligently if failed to take precautions against causing such harms commensurate with the reasonably foreseeable likelihood and magnitude of those harms. In other words, we expect people to take at least \$50 of precaution to avoid a one in ten chance of causing \$500 in harm to others.

The central idea behind the negligence model is that risk is unavoidable. Parties may be able to invest in mitigating risk, but cannot eliminate the possibility of risk entirely. If we were to hold parties responsible for any harms that they may cause to others, we are concerned that parties will either over-invest in precaution or avoid risky, but socially valuable, activity. For instance, driving is

inherently risky – on any given drive there is a chance that you will get into a costly accident. If we put too high a burden on drivers to avoid such accidents, they may over-invest in safety or avoid driving. But by only holding parties responsibly for taking *reasonable* precautions – that is, those commensurate with foreseeable harms – we don’t dissuade any socially-beneficial activity. In other words, modern negligence liability is designed to ensure parties engage in the socially optimal level of activities.

But negligence isn’t the only approach to assigning liability. Starting in the 1960s, courts began to impose so-called strict liability in some cases. Under a strict liability regime, a given party is always responsible for the harm incurred by its counterparties, no matter how careful that party was to avoid such harm. The underlying theory is that one party may be in a better position to prevent or assess likelihood of certain harms than the other. An important situation where this is the case is where assigning liability to one party allows for risk pooling, such as where parties on one side of a transaction systematically may not be able to absorb costs, or the expected costs may be distributed too thinly to justify taking precautions. It is also the case where one party is in a better informational position than the other, or is in a better position to gather or disseminate information.

Counterintuitively, as will be considered further in Part IV, strict liability does not affect the level of care that a party will take. One intuitively expects that if we impose strict liability on a party that that party will take greater precautions to avoid such harms than if it is only liable in the event of its own negligence. But this is not the case: under either model, parties will only invest in avoiding harms up to the point that the cost of such investment is commensurate with the expected magnitude and likelihood of harm. In other words, it never makes sense to spend \$75 to mitigate a one in ten chance of causing \$500 in harm. Rather, under strict liability, one will spend \$50 in precaution, and simply pay the balance of \$450 if that \$500 in harm happens to occur. The key difference between negligence and strict liability is how the \$450 loss that results when the harm does occur. Both negligence and strict liability accept that bad things happen – and that no amount of precaution can prevent every risk, and in fact that we do not want people to invest in inefficient levels of precaution. The difference between the two systems is that in a negligence regime the harmed party bears the cost of the harm; in a strict liability regime it is born by the other party. While the strict liability model may seem grossly unfair, we will see in Part IV that it turns out to function much like an insurance system and that it can, in fact, be a very positive model in some circumstances.

III.B. Cybersecurity is a Classic Case for Strict Liability

Cybersecurity presents a near textbook case for strict liability. The policy rationales for strict liability – the challenges that strict liability evolved to address – match the challenges created by cybersecurity. Indeed, even the historical challenges that gave rise to modern strict liability map onto the issues faced today in the cybersecurity setting. And while there are a number of common concerns about strict liability – concerns that militate against its use in various settings – they are largely inapposite to the cybersecurity setting.

The origins of modern strict liability in the American legal tradition are generally traced to Judge Cardozo’s famous opinion in *MacPherson v. Buick Motor Company*, which eliminated the privity

requirements for suits brought in tort. Prior to *MacPherson*, individuals could only bring suit against those with whom they shared some direct connection (that is, with whom they had privity). In other words, if a driver were injured in an automobile accident caused when a component of her car failed, she could only sue the person who sold her the car – she could not sue either the car manufacturer or a third-party manufacturer of the component that failed (for instance, if the component that failed had been bought by the manufacturer and integrated into the final component, as is often the case with many automotive components, such as tires). All of these parties only have an indirect relationship with the driver, so are said to lack privity. Under the pre-*MacPherson* model, the driver would be expected to sue to person who sold her the car, and that person could then countersue other third parties, either under separate legal theories or seeking indemnification.

MacPherson changed all of this, opening the door to direct suits by drivers (or other end-users) against manufacturers (or other responsible parties in the supply chain). The underlying rationales were intended to address the same sort of challenges that we see in the cybersecurity context. The multiplicity of parties in a supply chain make it difficult to figure out who to sue and make it difficult to establish or apportion liability. In both the case of *MacPherson* and the modern cybersecurity setting, this effectively externalizes risk onto consumers, and creates perverse incentives for how the various entities through the relevant product ecosystems design their products and services.

MacPherson was only the first step towards the modern understanding of strict liability. While it allowed parties to bring suits in the absence of privity, those suits were still brought under a negligence standard. Starting in the 1960s some courts began developing the modern understanding of strict liability in cases involving consumers harmed by (arguably) defectively designed or manufactured products. The canonical case is *Greenman v. Yuba Power Products*, which involved a wonderfully monstrous power tool sold by Yuba, the “Shopsmith, a combination power tool that could be used as a saw, drill, and wood lathe.” Mr Greenman was injured a year or so after his wife purchased Shopsmith for him as a Christmas present, and brought suit for breach of warranty and negligence. The trial court determine that the manufacturer had not been negligent, and that Mr. Greenman’s harms were not covered by any express or implied warranty.

On appeal, the Supreme Court of California found the manufacturer strictly liable for Mr. Greenman’s injuries, explaining that the question of “liability is not one governed by the law of contract warranties but by the law of strict liability in tort.” As explained by the court, “A manufacturer is strictly liable in tort when an article he places on the market, knowing that it is to be used without inspection for defects, proves to have a defect that causes injury to a human being. . . . The purpose of such liability is to insure that the costs of injuries resulting from defective products are borne by the manufacturers that put such products on the market rather than by the injured persons who are powerless to protect themselves.” Critically, under a strict liability model parties are not free to assign risk of harm by contract – any contract or warranty attempting to do so is a legal nullity.

Strict liability is an exception to the ordinary rule of negligence – it is only used in certain cases. The traditional examples are products liability, such as was the case in *Greenman*, and so-called ultrahazardous activities, such as keeping dangerous animals as pets or the use of explosives. It is clear why we only turn to strict liability in cases like these – and why cybersecurity is a similar case –

when we look to the core policy rationale underlying strict liability: ensuring that liability for harms be assigned to parties best able to bear it. Both negligence and strict liability accept that some amount of harm naturally occurs in the world. Under a negligence model, we assume that parties bear, and are able to bear, comparable responsibility for preventing or accepting the risk of harm. Under strict liability, we assume that the parties – especially in their abilities to prevent or accept risk – are asymmetric. In terms of *preventing* risk, we are generally concerned about risks that would be unreasonably, or impossibly, costly for individuals to detect. For instance, if a consumer could not determine whether a saw blade contained manufacturing defects without engaging in destructive testing of the blade, the law may hold the blade’s manufacturer strictly liable for manufacturing defects. Similarly with ultrahazardous activities, where it is unreasonable, for instance, to expect individuals to take precautions against the use of explosives in construction operations or pet tigers that may be roaming the streets – in these cases, too, we place a strict burden on the party engaging in the atypical activity. We see the same in terms of *accepting* risk. In this case, the concern relates to the parties’ relative abilities to bear the costs of a risk should harm come to pass. Again, the example of an injury related to power tools is illustrative: such an injury could be physically or economically devastating to an ordinary consumer, and many consumers do not have the knowledge or wherewithal to insure against such losses. The manufacturer, on the other hand, is in a much better position to assess the possible risks, and to insure consumers against those risks.

These rationale for strict scrutiny are not without criticism or nuance. Many of the criticisms of, and concerns raised by, strict liability regimes will be considered shortly below – and the idea that strict liability acts as a form of insurance will be considered in greater depth in Part IV.

Before considering these issues, we can outline the case for applying strict liability in the cybersecurity context. We already saw that the rationale for the first steps towards strict liability – *MacPherson’s* abandonment of privity requirements – mirrored concerns similar to those we see in the cybersecurity context: the difficulty of attributing liability and recovering damages that results from the multiplicity of actors and the complexity of their interconnected relationships. So, too, do the concerns about parties’ relative abilities to prevent and accept risk motivate modern principles of strict liability mirror reality of the cybersecurity setting. As discussed in Part I, every entity involved engaged in conduct online – from individuals, to small businesses, to non-profit and governmental organizations, to large non-tech firms, to large tech firms – is exposed to cybersecurity risks. Mitigating these risks is far beyond the expertise of the vast majority of these entities. And, even if it weren’t beyond their competence, most defects in third-party systems are latent. Even if these third parties open their systems up for inspection, it is functionally impossible to expect even sophisticated parties to audit them for defects at reasonable costs.

This is largely descriptive of the situation that exists in the business-to-business landscape. Even among sophisticated parties, few are in a position to meaningfully understand, let alone prevent, cybersecurity risks. But this is even more dramatically the case in the consumer-to-business relationship. Here, consumers are almost entirely at the mercy of the firms they interact with online to keep data that they disclose to those firms secure. Consumers have no visibility into those firms’ systems, into what data those firms retain, how they manage that data or use it, or to whom it is disclosed (intentionally or unintentionally). Once their data has been shared with a firm, consumers

have literally no ability to monitor its subsequent use or handling, to take precautions to prevent harm, to detect its misuse, or to take action in response to those harms.

Indeed, the cybersecurity context arguably presents a more “textbook case” for the use of strict liability than seen in most “textbook cases.” Strict liability regimes have two basic effects: they increase the price of products and services, and they encourage risk-taking by consumers. They increase the cost faced by providers of products and services because they those providers bear the risk of any liability. But these costs are almost always passed on to consumers. The net effect, discussed in more detail in Part IV, is that firms subject to strict liability act as insurers: they spread the cost of risk across the entire pool of consumers, collecting a premium for that risk through the price they charge, and they use those premiums to pay out damages as they occur on a stochastic basis. Importantly, the concern about increases prices has an important, potentially pernicious, secondary effect: more price sensitive consumers, or those who are less exposed to risk, may select themselves out of the market. This has the effect of spreading the cost of risk across smaller pool of consumers, each of whom therefore has to pay proportionally more. Taken to the extreme, this can make some products unviable in the market – as was indeed observed in the 1980s. In proposing a strict liability regime, we need to be very cautious about this concern, as it could be devastating to the market. The second concern is similarly important: if consumers are aware (implicitly or explicitly) that another party is liable for any harm that befalls them, consumers may have an incentive to opportunistically engage in riskier behavior. For instance, consumers may shirk on routine maintenance of potentially dangerous products, or fail to read manuals or otherwise educate themselves to the safe operation of potentially dangerous products, if they know that they will receive compensation despite their own negligence.

There are two well-understood problems raised by strict liability: adverse selection and moral hazard. Fortunately, neither of these concerns is substantial in the cybersecurity context. Users and purchasers of products and services throughout the cyber-domain consistently have little ability to control, monitor, or prevent against harm. At the retail level, consumers are wholly at the mercy of the firms with which they work and to which they provide data to ensure that that data is reasonably stored, used, and secured. The best that even a sophisticated consumer can do is rely on a firm’s assurances and reputation. But if the past several years have demonstrated anything about security, it is that even security-conscious, sophisticated firms can be the subject of cyber-incidents. The same also holds in the business-to-business context. When one firm engages another to provide security-related services or products, it is generally because the contracting firm lacks the sophistication or resources to implement those products or services on its own. And, as discussed at the beginning of this Article, the complexity of software and of designing secure systems, means that contracting parties cannot reasonably audit or monitor the performance of most security-related products or services.

Taken together, this analysis means that neither of the common concerns about strict liability regimes apply in this context. A transition to strict liability likely *will* increase the cost of providing products or services, especially in the short run, *but not* in a way that is likely to adversely affect consumers. The risk of harm from cyber-incidents is spread relatively uniformly across the online ecosystem, which means that we are not worried about price-increases causing some portion of the market to opt-out of the market (leading to further increases in price to the remaining

portion). Indeed, the opposite is more likely to occur: concerns about security today increase the cost of participating in these markets today, which may cause risk-averse users to opt out of the market (or to engage in costly and largely ineffective self-help, such as using complex password management systems or multiple e-mail addresses). Pushing the cost of these risks back to the parties best able to mitigate and bear them could actually grow the market, rather than segmenting it. And, in the long-run, placing the risk of cyber-incidents on parties that are better able to mitigate them will likely lead to an overall improvement in the systems that make up the cybersecurity ecosystem, reducing the overall risk for everyone. Similarly, the second concern about strict liability – that, in this context, it creates perverse incentives for users and contracting parties to engage in riskier behavior – is largely inapposite. Today, it is hard to imagine an environment in which participants routinely engage in riskier behavior.

The basic problem in the security ecosystem as it exists today is that the difficulty of imposing liability in negligence and contract models has effectively created a “strict fault” regime. Under this current regime – which is governed by negligence and contract law in name only – sophisticated parties pervasively externalize risk upon unsophisticated parties. This is exactly the opposite of how the law usually works, and of how we should want to see incentives structured: we generally want to impose liability in the first instance on the parties best able to prevent harm from occurring or to absorb the risk of harms that do come to pass. Doing so tends to reduce overall risk to society by maximizing incentives to efficiently reduce it and minimizing the costs of dealing with it. Under a negligence and contract model, we have seen the opposite: incentives to burden unsophisticated parties with risk rather than working to mitigate it, without any concern for the cost of the resulting harms. Strict liability is manifestly a better approach.

III.C. Limitations: Statutory Damages, and Other Practicalities, and Best Laid Plans

Transitioning to a strict liability regime would address many of the problems facing today’s cybersecurity ecosystem. Indeed, the most important advantages of a strict liability regime – that it serves as a form of insurance – won’t be addressed until the next Part of this Article. But strict liability is not a panacea. Before turning to discussing the relationship between strict liability and insurance, some practicalities of implementing a strict liability regime for cyber-incidents need to be considered.

The substantial limitation on a strict liability regime is that it does nothing to address the question of damages. Recall that one of the greatest obstacles to imposing civil liability on firms that have experienced data breaches, or other cyber-incidents, has been proving cognizable harm in court. Courts have consistently found that damages cannot be awarded in these cases because causation is too tenuous, there are too many potential intervening factors that could have caused any harm, or harm is too speculative to quantify.

There is a straightforward solution to this problem: statutorily-directed damages. “Statutorily-directed” means two things. First, courts should be instructed to err on the side of finding cognizable damages. Evidence still needs to be required to support a finding of damages, but courts can be statutorily directed to require a reduced burden of proof, shift the burden of proof, or accept that certain harms (e.g., privacy harm) are cognizable. Second, and more important, Congress

can direct the establishment of a schedule of damages to be used by courts in establishing damages for various sorts of harms at trial. The Federal Trade Commission (FTC), for instance, could be directed to establish such a schedule of damages through a rule-making process, with instruction that the schedule be based on empirical data but that the agency should err on the side of finding substantial damages and that deference should be given to the agency in interpreting that data. Indeed, for reasons discussed in Part IV, it would make sense to require the agency to use a multiplier in setting damages. Once such a schedule of damages had been set, courts would use it as a floor in the civil context – a judge could still find higher actual damages or assess punitive damages where appropriate – and would otherwise fall back on the statutory direction to find cognizable damages in the event of harms not covered in the schedule.

[[[Add discussion of *who* is subject to SL regime to below paragraph. Those who collect user/3rd party data, who sell/distribute/provide systems and services that store/manage or control access to such data, and who integrate such systems are strictly liable to direct and indirect retail consumers, but can sue each other for negligence/indemnification. For instance, a systems integrator or hardware manufacturer could be strictly liable to a consumer or consumer-facing firm, even if it was a software fault that enabled the cyber-incident – *but* they could in turn sue the software developer for the defect in the software. And a consumer could sue a consumer-facing firm that collected her data and experienced a breach.]]]

Beyond the question of damages, there are other implementation details and decisions that would need to be addressed. A few specific points are discussed below, with the goal of designing a system that is broadly incentive-aligned – though it is certainly not the only approach, and there certainly are other issues that may need to be addressed. One challenge that the switch to strict liability does not address is the incentive that firms face to detect, disclose, and otherwise respond to adverse cyber-incidents. Simply stated, under any liability regime a firm will face no liability if it can keep an incident secret. As an initial matter, a federal civil cause of action should be created alongside the transition to strict liability that allows both private parties (acting alone or as a class) and the FTC to bring civil actions in federal court. The low bar to recovery created by the strict liability nature of this cause of action, along with the multiplier to be used by the FTC (or other agency) in developing a schedule of damages, creates an initial incentive for those potentially harmed by cyber-incidents to be vigilant in monitoring and taking action in response to them. Additionally, punitive damages should be expressly authorized – even encouraged – for firms that do not timely detect or respond to a cyber-incident. On the other hand, firms should be affirmatively encouraged to put procedures in place for the timely detection of and response to cyber-incidents – including providing notice and reasonable compensation to harmed parties. One simple approach to creating such an incentive is to bar suits by the FTC or class actions against firms that have such procedures in place.

IV. Data-Breach and Cyber-Incident Insurance: The Shield

Part III of this Article argued for the use of strict liability in addressing harms that result from cyber incidents. Use of strict liability in this context would correct many of the challenges that parties face in establishing liability for harmful conduct discussed in Part II. In particular, the failures

of existing private and public law mechanisms to assign liability for cyber-incidents creates incentives for those who would otherwise bear the cost of mitigating or the costs of harms caused by such incidents – generally the same parties who are in the best position to take precautions against them – to externalize the risk of cyber-incidents on to third parties – generally those least able to mitigate or afford to bear such risk. In effect, the current model is a no-liability model, which creates pervasive, harmful incentives. Adopting a strict liability model would go far to align public and private incentives to reduce cyber-incident risks to a more efficient level.

But there is another, more powerful, benefit to using strict liability in this setting: strict liability is, effectively, a form of insurance, and broad adoption of strict liability would foster the development of a vibrant market for data-breach and other cyber-incident insurance policies. An insurance-based approach to addressing cybersecurity concerns is substantially preferable to either the current strict-fault approach or even an effective liability-based system that focuses on individual actors. This is because cybersecurity-related risks are based in systemic problems that affect users throughout the cybersecurity ecosystem – they rarely result from individual-scale problems. Such risks are better managed on a pooled, insurance-scale basis. And, importantly, insurance-based systems have powerful regulatory characteristics that can both improve the behavior of individuals operating in the current cybersecurity ecosystem and improve the overall state of that ecosystem moving forward.

[[[To this end, this Part proposes ...]]]

IV.A. Strict Liability as Insurance

The intuitive understanding of strict liability is that it is meant to place the burden of avoiding harm on the more sophisticated party in a relationship – generally the party with greater knowledge about the risks associated with the use of a given product or service. The traditional example is of a dangerous product such as a power tool, which may have some latent defect or require non-obvious training in order to use safely. Ordinary consumers reasonably have need to use such products, but are frequently ill-equipped to do so safely. Strict liability seemingly places the burden of ensuring the safety of such products on the manufacturer – the sophisticated party – to ensure that the manufacturer goes beyond the requirements of ordinary negligence to also ensure that the product is safe for consumers who may themselves be negligent. This seemingly-reasonable understanding is explained on the grounds that the sophisticated party is in a better position to mitigate such harm than its counterparties, so the burden should be placed on the sophisticated party. As Justice Traynor, the author of the *Greenman* opinion commonly heralded as establishing modern strict liability, wrote in an earlier (dissenting) opinion, the one in which he first articulated his concept of strict liability: “public policy demands that responsibility be fixed wherever it will most effectively reduce the hazards to life and health inherent in defective products that reach the market.”

But this is not, in fact, how strict liability operates. Following early adoption of modern forms of strict liability, legal scholars realized – both through theoretical and empirical research – that strict liability does not induce firms (or other parties subject to strict liability) to take greater care to avoid harm than ordinary negligence. The reason for this is simple: ordinary negligence

encourages firms to take precautions commensurate with the expected likelihood and magnitude of harm. In other words, under ordinary negligence, firms will invest up to \$50 to avoid a 1-in-10 likelihood a \$500 harm (that is, an expected \$50 in harm). Counterintuitively, however, the transition to strict liability does not change this: a firm will not invest \$60 to prevent an expected \$50 in harm. Rather, it makes more economic sense for a firm to invest \$50 in precaution, hope that the harm does not come to pass, and, if it does come to pass, write a check to the harmed party.

Strict liability, in other words, does not increase either party's incentives to take precautions against harm. Indeed, unless it is implemented with a contributory-negligence defense, it can *encourage* negligent behavior on the part of the non-liable party, since that party knows it is effectively insured against harm by the strictly-liable party. Rather than affect parties' incentives, strict liability's real effect is to shift risk of harm from one party to another.

The practical consequence of shifting risk in this way is that strictly-liable parties become insurers for their counterparties. This is most easily seen in the case of strict products liability. In the first instance, firms will invest in precautions – such as designing their products to minimize the risk of harm through ordinary use – up to that point where the investment in precautions equals the expected likelihood and magnitude of harm. Beyond investing in cost-effective precautions against foreseeable harms, firms will also set aside a portion of their revenues to cover the cost of foreseeable harms that cannot be cost-effectively mitigated. Thus, a firm will invest \$50 in precaution to avoid a 1-in-10 chance of a \$500 harm (or, more precisely, to minimize the expected likelihood and magnitude of harm as a function of the incremental cost of investment in precautions), but it will also set aside \$50 for each unit that it sells in order to compensate the one in ten consumers whom harm is expected to befall. Critically, that \$50 set aside for each unit sold does not come from the manufacturer. Instead, the manufacturer passes those costs along to the consumer, increasing the price of its products in order to meet the strict liability regime's demand that it insure consumers against harm.

In the products liability market, this is, literally, exactly how strict liability works: firms purchase third-party insurance for the users of their products and incorporate the cost of this insurance into their products' prices.

In effect, adopting a strict liability regime is equivalent to adopting a mandate that parties have insurance. Such a mandate makes sense in many contexts. In the products liability context, for instance, manufacturers have much greater ability to inspect their products and detect latent defects, or to provide basic instruction on the safe use of their products. Even more important, however, they are in a better position to understand the risks of using their products and provide (or purchase) insurance against those risks. It would be very costly to expect the relatively large number of consumers in the economy to research and purchase separate insurance policies for each product they happen to buy – indeed, in most cases more costly than the value of the product to the consumer. On the other hand, it is relatively inexpensive to require a relatively small number of manufacturers to purchase third-party policies. And, on the insurer side, it is relatively costly to negotiate individual policies for individual consumers, each of whom has highly individualized characteristics – but it is relatively inexpensive to negotiate a third-party policy for a large pool of consumers.

The same analysis holds in the cybersecurity context. As was discussed in Part III.B, the difficulties of assigning and quantifying risk under the negligence regime has allowed relatively sophisticated parties to systematically externalize risk onto relatively unsophisticated parties. Under today's model, sophisticated parties underinvest in security and impose the cost of security risks on to unsophisticated parties. The first order effects of a transition to a strict liability model would be to encourage investment in precautions against cybersecurity-related risks. But given the relatively large number of relatively diverse and unsophisticated users engaging in diverse online activities that expose them hard-to-understand and hard-to-quantify risks – risks against which users are often not only ill-equipped but entirely unable to mitigate – as compared to the relatively small number of relatively sophisticated firms on the other sides of these interactions – firms that often entirely control all of the security-relevant aspects of their relationship with consumers (e.g., how and what data use stored and used, and the systems that control access to that data) – it makes very good sense to require those providing security-related products and services to insure users of those products and services. And, to be clear, “security-related products and services” includes security-related services bundled with non-security related products and services. This would include, for instance, firms that collect and store consumer data implicitly bundle systems for the secure storage and use of that data. And it would also include, for instance, the sale of products that can reasonably be expected to be used to manage a firms security-related systems, such as server software or other software that allows for storage of and access to data.

Of course, just as the risk-mitigation understanding of strict liability was explained to be “no panacea” in Part III, it also is no panacea under the insurance understanding. Indeed, and unsurprisingly once we recognize that strict liability is really an insurance regime instead of a liability regime, the same analysis offered in Part III applies here, as well. The two basic concerns about strict liability – that it will increase costs, possibly pushing some participants out of the market and thereby further increasing costs for those who stay in, and that it creates perverse incentives for parties to engage in riskier behavior – are in fact the same two basic concerns inherent in the insurance context. In the terminology of insurance, the first concern is called “adverse selection” – that the cost of insurance will cause lower-risk parties to drop out of the market, leaving only the higher-risk, higher-cost, parties behind, which in turn will further increase the cost of insurance. And the second concern, that of insurance create perverse incentives to engage in riskier conduct, is called “moral hazard.” But for the same reasons discussed in Part III, these do not provide a persuasive argument against adopting strict liability, including its insurance-like characteristics, in the cybersecurity context. Risk is pervasive throughout the cybersecurity ecosystem, such that there is little concern that a meaningful portion of the ecosystem will adversely select-out of the market due to the potential cost of insurance – indeed, correcting the existing perverse incentives that sophisticated parties face to shirk on security (which, to be clear, would be corrected by the liability aspect of strict liability, not the insurance aspect), is likely to drive further participation in the market. And the insecurity of the current online ecosystem is so great, and participating in that ecosystem already so risky, that it is hard to imagine any concern about moral hazard.

IV.B. Insurance as Regulation

The discussion so far has focused on the effects of strict liability on individual actors, or in the context of individual relationships – how liability and risk is apportioned between individual users of security-related products and services and those providing those products and services. The insurance-like characteristics of strict-liability, however, offer another powerful argument for the application of strict liability in the cybersecurity context: insurance can have a broad regulatory effect on the overall market. A likely outcome of a transition to a strict-liability regime for cybersecurity is that most firms – especially relatively unsophisticated firms – would purchase third-party cyber-incident insurance policies. This alone would be the single most important step to improving the overall quality of the cybersecurity ecosystem that is possible today, because insurers have unique abilities and incentives to inform firms of their legal obligations, to develop best-practices, to audit and educate firms as to those practices, and to lobby for improvements to the overall state of the cybersecurity ecosystem.

Part I of this article made the case that, stated simply, cybersecurity is hard. Participants throughout the cybersecurity ecosystem have surprisingly little understanding of the pervasive risks and challenges associated with cybersecurity, let alone an understanding of best (or even merely good) practices. At the user-level, individuals rarely have familiarity with principles of security “hygiene,” as it is referred to in the literature – concepts like how to identify and respond to risks, how to manage passwords, e-mail accounts, and other sensitive information, who to (and not to) share information with. Going up a level in the food chain, small businesses frequently lack all of this knowledge, as well, but are also tasked with bigger, more complicated challenges: how to design and secure basic IT systems, how to monitor, identify, mitigate, and respond to cyber-incidents, how to design and implement an incident response plan. For most small- and even mid-sized businesses, it is enough of a challenge (and expense) just to get a basic IT infrastructure in place. Especially outside of the tech sector most firms don’t even know what issues they face, let alone how to address them. The same basic story can be told at nearly every level, and for nearly every level of sophistication, of participation in the cybersecurity ecosystem: even those implementing systems often lack information about the needs and sophistication of their customers, making it difficult to design systems that promote good security hygiene.

Insurance – and insurers – are uniquely situated to address these system and ecosystem-level concerns. The basic business of insurance is the collection of data relating to, and quantification of, risk, the evaluation of individual insured’s risk profiles, and the minimization of both their insureds’ exposure to and the overall market’s creation of risk.

An insurer’s first job is to build its actuarial tables – to collect data that will allow it to evaluate individual insured’s exposure to risk in a given industry. This is a process that, in the context of cybersecurity, we will return to in Part IV.C.

Once an insurer has a sense of the factors that go into determining an individual insured’s exposure to risk, it can begin underwriting new clients. This is the process by which insurers evaluate possible clients’ risk profiles to determine their insurability and the premiums to be charged for insurance. In the cybersecurity context, for instance, an insurer is likely to conduct an audit of a firm’s systems and procedures: what is the architecture of the firm’s network, what data is stored

and used, how is access controlled, what incident response plans are in place, how are employees trained and monitored, how is third-party access to the firm's systems (e.g., by contractors) assigned and monitored, and the like. This process alone offers – or would offer – most firms a more in-depth evaluation of their cybersecurity systems than they would ever otherwise receive, except possibly in the case of a serious security incident. Even more important, it would offer these firms the most in-depth education regarding security best-practices that they are likely to ever receive. Moreover, insurers have an incentive to monitor their insured's ongoing performance, providing ongoing updates and training regarding best practices and responses to newly discovered security problems

It is difficult to imagine a more effective approach to educating and evaluating the bottommost layers of the security pyramid. And this leads to a second powerful benefit of widespread adoption of cyber-incident insurance policies: effective push-back against the pervasive externalization of risks from sophisticated parties onto unsophisticated parties. There are two sources of this push-back. First, as consumers and firms are informed by insurers about their exposure to risk, they will have greater demand for more secure products. This means both that they will be willing to pay more for more secure products (which would lower their insurance premiums), and they will put greater pressure on vendors and service providers to provide more secure products and services. And second, the insurance industry itself will serve as a powerful lobby to push for better designed, more secure, products and services. The software and tech industries are powerful interests that have largely been successful in shielding themselves from liability for the quality of their wares – for good reasons and bad. There is no concerted interest group on the other side of this balance – consumers are too diffuse a group with too little an understanding of the relevant harms to which they are exposed to effectively lobby against Silicon Valley or the Business Software Alliance to demand risk be shifted back to those who have built the insecure infrastructure on which we have come to depend.

IV.C. Jumpstarting the Market

There are two approaches that a firm subject to strict liability can take to insuring against harms for which it will be held liable: it can either self-insure or purchase third-party insurance. Both approaches are commonly seen in the products liability context, and there are specific circumstances under which either approach may make sense. In the case of cybersecurity, however, it will generally make sense for all but the largest firms to turn to third-party insurance. This is simply because most firms don't have the necessary data or sophistication to construct the actuarial data needed to evaluate their customers' risk exposure, to apportion the cost of that exposure across their customer base, or to reasonably manage claims. Rather, it will generally make much more sense for firms to purchase third-party liability coverage from an insurer with expertise in cybersecurity, with data aggregated from a across a range of sources and security incidents, and with the resources necessary to manage the sort of claims seen in the cybersecurity context.

This poses a substantial problem: today there is no robust market for cyberinsurance, especially for small- and mid-size firms. Insurers have been writing cyberinsurance policies – many of which contain cybersecurity and related coverage – for large, and relatively sophisticated, firms

for a number of years. But these policies are generally bespoke, written for specific firms and based upon those firms' individualized needs and characteristics. Such policies are inappropriate to the vast majority of parties potentially subject to strict liability for adverse cyber-incidents. The transition to strict liability needs to be accompanied by the creation of a retail marketplace for cybersecurity and cyberincident insurance.

[[[Unfortunately, there are a number of obstacles to the development of such a market. Limited actuarial data; correlated risks, not iid; regulatory uncertainty – unclear what firms' legal exposure to these risks are (due, e.g., to FTC efforts, legislative uncertainty).]]]

Fortunately, the transition to a strict liability regime described in Part III contains within it a mechanism for jumpstarting such a market: the schedule of statutory damages. The previous discussion of statutory damages indicated that they should be set and interpreted aggressively – that the FTC (or other agency) and courts should be instructed to err on the side of finding that damages are not only cognizable, but substantial. Indeed, it makes sense that these statutory damages should be set at a multiple of estimated actual damages. This is in part to correct for infrequency of detection. But it is also in part to create an incentive for the development of a robust insurance market.

In order to “jump start” the market for retail-scale cyber-incident insurance, damages paid by insurers would be limited to actual damages, as provable by insurers' actuarial data. The schedule of statutory damages would be taken as presumptively valid, but insurers would be able to rebut them as actual data about the costs of various types of cyber-incidents was collected. This would have two effects. First, the relatively high statutory damages would create demand for insurance products by firms exposed to cyber-incident risks. And, second, insurers would be able to charge relatively high premiums as this market develops – in order for their products to be desirable, they would only need to charge premiums incrementally below the potential costs firms face from aggressive statutory damages. As insurers became more familiar with the market and develop actuarial data, the cost of these products will then be driven down as firms compete for business.

Summary and Conclusion