Chesney on Cybersecurity Law, Policy, and Institutions

v.3.0 (March 2020)

Professor Bobby Chesney
The University of Texas at Austin
@bobbychesney

About this book: This is an interdisciplinary "eCasebook," designed from the ground up to reflect the intertwined nature of the legal and policy questions associated with cybersecurity. My aim is to help the reader understand the nature and functions of the various government and private-sector actors associated with cybersecurity in the United States, the policy goals they pursue, the issues and challenges they face, and the legal environment in which all of this takes place. The first part of the book focuses on the "defensive" perspective (meaning that we will assume an overarching policy goal of minimizing unauthorized access to or disruption of computer systems). The second part focuses on the "offensive" perspective (meaning that there are contexts in which unauthorized access or disruption might actually be desirable as a matter of policy). In short, the book is a guided tour of the broad cybersecurity landscape.

Who should use it? The book is designed to be valuable not just to beginners but also those who may have experience in one area but would like to see how their corner of the puzzle relates to the larger whole. At the University of Texas, I use it as the main text in an interdisciplinary course that includes students from our schools of law, public affairs, computer science, engineering, information, communications, and business. My aim is for it to work equally well, however, as a free-standing text for those who want to study these topics independently.

Why did you write this book and why is it free? I wrote this book under the auspices of the interdisciplinary cybersecurity program at UT's Robert Strauss Center for International Security and Law, with generous support from the Hewlett Foundation's Cyber Initiative.[1] We believe that progress in this area has been inhibited to an unnecessary degree by a lack of mutual understanding among lawyers, engineers, computer scientists, government officials, and business leaders, and we have set out to develop courses that will help close that gap. Wanting to ensure the widest possible distribution and impact in light of this goal, and to have the option to quickly and easily update the text when needed, I opted to make the book freely available online rather than selling it through a publisher. I hope you'll find the end result appealing, even if it is much less-formal and polished than would have been the case otherwise. And, if you do, I hope you will further the mission by recommending the book to others.

[1] I extend special thanks to Rachael Jensen for her outstanding research and teaching assistance, especially her key role in adapting earlier versions of these readings into the current format, and to Eli Sugarman of the Hewlett Foundation.

Table of Contents

Note: This book is broken up into 25 separately-numbered subsections in order to make it easy to use in a 3-credit course, with each item scaled to suit a 60-to-90 minute class session.

I. THE DEFENSIVE PERSPECTIVE

For the first part of this book, we will focus on the *defensive* perspective. That is, we will focus on the overarching public-policy goals of (i) minimizing unauthorized access, disruption, manipulation, and damage to computers and (ii) mitigating the harm when such malicious activity nonetheless occurs.

When it comes to the first of those two core goals, the main idea is simple: we want to increase the undesirable consequences attackers believe they will risk should they attempt certain actions, while also making it more difficult for them to succeed. Part A below focuses of the former, and Part B below focuses on the latter. In both contexts, our goal is to understand the various institutions, policies, and legal frameworks that define the status quo on these matters; to grasp the competing interests that are in play; and to wrestle with the question of how we might do better. Part C then takes up the second major goal noted above: managing the consequences of successful attacks.

Note: At this point, you might be thinking something along these lines: sure, deterrence and resilience are fine as far as they go, but what about achieving defensive goals by using cyber capabilities directly to disrupt the capabilities an adversary would need in order to attack in the first place? Isn't that "defense" too? Yes, at least from a certain point of view. And we will indeed discuss "disruption" as a strategy distinct from deterrence in this unit. But we will do so only at a conceptual level; we will save our detailed study of the institutions, legal frameworks, and policy challenges associated with this approach for the second part of the book ("II. The Offensive Perspective"), on the theory that such operations may be defensively-motivated yet nonetheless involve unauthorized out-of-network activity best examined in context with other activity of that kind.

A. Imposing Costs on Attackers

In this section we will study an array of tools that the government can use to impose costs on attackers, starting with the tools of criminal law enforcement and then moving on to civil liability and beyond that to an array of options that become relevant when the attacker is or might be associated with a foreign government.

1. Introduction to Key Terms and Concepts

We will begin by establishing common ground regarding key concepts and terms associated with cybersecurity.

A. About those "attackers"

In this book we often will use the word "attacker" as a shorthand referring to a person or organization that seeks to access, disrupt, manipulate, or damage a system in an unauthorized way. Attackers come in many shapes and sizes. Some are sophisticated professionals, others are rank amateurs. Some are state-sponsored, some are part of non-state organizations (and some of these nonetheless sometimes work on behalf of states), and some are individuals. Some are crooks. Some are spies. Some are just showing off skills. Some are in it for the laughs. Some do it to settle personal scores. Some are seeking competitive advantage. Some mean well, hoping to spur people to increase their defenses by exposing weaknesses in hopes that they'll be remedied. Some are malicious, hoping to cause harm (or to use your system to cause harm to others). The point being, there are *many* potential attackers out there, with a wide variety of motives and

capacities, some awful and others laudable. Bear this in mind as we examine the various tools that we currently have—or might one day have—to impose consequences on attackers.

B. Core Goals Associated with the Security of Information

There is a particular organizing principle often used to explain the core goals of information security ("infosec"): the "CIA Triad." That's not a reference to the Central Intelligence Agency. Rather, the acronym refers to three distinct defensive policy goals: preserving the confidentiality, integrity, and availability of information.

Confidentiality:

> You may sometimes be happy to expose certain information to the world, but often you'll want to control access to it. Ensuring the confidentiality of information means just that—ensuring that people can access it only in accordance with the terms the owner of the information sets. We can also refer to this as privacy, though "confidentiality" fits a bit better and is more commonly used in this setting. Those who gain unauthorized access to an information system can cause harm, from this point of view, by violating confidentiality—either by just reading the information or exfiltrating a copy of it out of the system.

Integrity:

> In addition to controlling the confidentiality of information, we also commonly wish to ensure that information is not altered or destroyed in an unauthorized way. Those who gain unauthorized access to a system at a level permitting alteration of data are thus in a position to cause harm to the integrity of the information.

Availability:

> Another goal for the security of information systems is to ensure systems work as intended—that is, that the information remains available for its intended use. An attacker who cannot pierce the confidentiality of information or harm its integrity might nonetheless be able to disrupt access to it, causing harm to the availability of the information. And notice, too, that measures one might take to protect against harms to the confidentiality and integrity of a might impose undue costs along the availability dimension.

In short, the overarching goal of cybersecurity from the defensive perspective is to prevent (or, at least, minimize) unauthorized, harmful actions along any of these dimensions.

C. More Key Terms

Vulnerability:

> In *any* system, there are latent weaknesses. Those weaknesses present opportunities that might be exploited purposefully (or even just accidentally) to cause the system to perform in a way its owner does not intend. This is endemic to software. The code that comprises software tends to be complex, sometimes mind-bogglingly so. Vulnerabilities—also known as "bugs" or "vulns"—are bound to be there. Some are easily spotted, and some require sheer genius or blind luck to discover. Some are easily remedied (that is, "**patched**" by a change to the code in question), but some cannot readily be fixed without disrupting the functionality of the code.

Note: Just because a vuln has been identified does not mean the creator/supplier of the software at issue will develop a patch. And even if a patch *is* developed, it does not follow that everyone using that software will apply the patch. Many vulns remain useful to attackers long, long after their existence is discovered and a patch developed.

Exploit:

Just knowing that vuln exists is not enough to provide someone with unauthorized access to a system—some further step is required to take advantage of the vuln. An "exploit" is a program or technique that takes advantage of a vuln in order to achieve some unauthorized effect within the system.

Note: An attacker may not immediately have access to all parts of the system at issue. "**Privilege escalation**" refers to additional steps that extend the attacker's unauthorized access. An "**exploit chain**" refers to the combination of programs and techniques that an attacker assembles to gain access to a system, escalate privileges within it, and then perhaps to take further actions.

Disclosure:

A person who discovers a vulnerability (or, for that matter, who identifies an exploit targeting a vulnerability) might choose to disclose the information to some responsible party, such as the creator of the software that suffers from the vulnerability or to a public database of known vulns. Some companies actually pay bounties for such information through "**bug bounty programs**," but then again, some companies have unfortunately suspicious (or worse) reactions to receiving such information out of the blue—especially if they perceive the situation as an extortion attempt. Meanwhile, not everyone who discovers a vuln is inclined to disclose it. Some would prefer to sell the information on the black market, or even keep it for their own eventual use.

A "**zero-day vulnerability**" is a vulnerability that has not been disclosed. The name reflects the fact that is has been zero days since disclosure, and thus there is not yet a patch for the vuln (let alone distribution of a patch). Many assume that this makes zero-day vulns the most valuable bugs. But, in fact, the value of a vuln depends on many factors, especially the nature of the access that follows when the vuln is exploited successfully. There are many **x-day vulns** (i.e., vulns that were disclosed x days ago) that have continuing value, and some zero-days that are not particularly valuable.

Note: The fact that a company receives notification of a bug does not mean the company will necessarily work to develop a patch for it. Most companies have a primary mission (selling products or services) where they would prefer to spend their resources. Investing the time, money, and labor in increasing cybersecurity diverts resources away from profit-seeking projects. And so, when a company is alerted to a vulnerability in its system, it will undoubtedly weigh the costs and benefits associated with patching it.

Social Engineering:

Alas, the greatest and most persistent vulnerability is, well, us. "Social engineering" means tricking people into revealing information or otherwise taking an action that make it possible for someone to gain unauthorized access to a system. This can be as simple as tricking someone into sharing a password, or as complicated as developing a highly customized email designed to get the target to open an attachment that contains malware.

"**Phishing**" is the generic term for tricking someone into opening an attachment or clicking a link that opens the door for malware to access a system. "**Spear phishing**" is when the inducement is tailored to the target—for example, an email that appears to be from a coworker asking you to review an attached document.

Another category of human vulnerability involves "**insider threats**"—that is, people within an organization who properly have access to a system but then use that access in an unauthorized way or cooperate with someone else to enable them to do so. Some insider threats are themselves malicious, such as an employee stealing trade secrets from his employer. Others are induced by an outside actor working through an insider—whether the insider knows it or not. For example, an employee who forwards an email that contains links or attachments laced with malware is an insider who poses a threat to his organization, even though he is unaware. And don't disregard witting forms of collaboration resulting from bribery, blackmail, coercion, etc.

Spend some time thinking about the various reasons why someone might want to gain unauthorized access to a system, and what types of individuals and organizations might act on those reasons. Be prepared to offer your thoughts during class, as we work collectively to map out the constellation of motives and actors.

D. Meet the Black Market

When a person or entity wants to gain unauthorized access to, or disrupt, a computer or network, it certainly helps if that person or entity already has the skill and resources needed to develop tools to suit that purpose. But not everyone does, and for the most part, they don't actually need to. Why not? Because there is a thriving black market for the sale of both stolen information (credit card numbers, etc.), and the tools needed to steal and disrupt.

Read chapters 2–4 of this 2014 RAND Study, and consider the following questions:

- Who participates in these black markets, and has the answer to that question changed over time? Note factors such as nationality, and expertise. What are the policy implications of your answers?
- How have the types of products/services sold on the black market changed over time? Why does this matter from a policy perspective?
- Botnets: What are botnets, and how has their use evolved over time? What is the relationship between the Internet of Things (IoT) and the botnet problem?
- Sometimes government does succeed in "taking down" a particular dark web market. But the RAND report suggests such successes are "transitory." Why would that be, and what policy insights follow from this?

Obviously, anonymity is important to participants in cybercrime black markets. Read this January 2017 Wired article from Andy Greenberg for an introduction to how these markets and their participants try to remain hidden.

As you read, consider the following questions and prepare to discuss them in class:

- Define the terms "deep web," "dark web," and "TOR."
- Which policy arguments might favor allowing at least some such hidden services to exist? Which favor suppressing them?
- How should these interests be reconciled? Should the balance should be the same in all societies? Why might that be hard in practice?

2. The Crime Model: Key Institutions and the CFAA

In this unit we are assuming the paramount policy goal is to minimize unauthorized access to (or disruption of) computers, and we open by examining tools that advance this aim by imposing costs on attackers. Among the most visible and discussed of these tools is criminal prosecution.

A. Meet the Government's Investigators

There are a surprising number of government agencies responsible for investigating cybercrime, and their roles often overlap. As you read the information below, you might want to create a list or chart that identifies each agency, the relevant department/structure within the agency, and their basic responsibilities. Throughout the semester, it may be useful to refer back to this list and update as needed.

Department of Justice

As you read here and here to learn about the Department of Justice's role in investigating and prosecuting attackers, consider the following questions:

- What is the office at DOJ that has special responsibility for this area, in general?
- Which office at the Justice Department appears to have had lead responsibility in the APT 10 Group case, and why do you suppose that is?

FBI

As you read here and here to learn about the FBI's role in investigating and prosecuting attackers, consider the following questions:

- What part of FBI focuses on computer crime, in general?
- Can you specify the various roles FBI plays?

U.S. Secret Service

As you read here about the Secret Service's role in investigating and prosecuting attackers, consider the following questions:

- What role does the Secret Service play?
- Why is the Secret Service involved in this area?

We are focused for the moment on prosecuting hacking as a crime. But it's worth pausing to remind ourselves that "crime" is not always the most relevant label to describe an unauthorized access scenario, even if it does involve a violation of criminal law.

So stretch your mind a bit, recalling the reference in the readings above to DOJ's National Security Division:

- When DOJ decides to prosecute, how might this have implications—perhaps negative ones—for the missions of other U.S. government agencies or departments?

B. The Computer Fraud and Abuse Act

There are *many* federal crimes that might be implicated by the activity we are discussing, but the most significant one is the Computer Fraud and Abuse Act, or CFAA. The CFAA is codified in Title 18 of the United States Code (the U.S. Code is the compilation of federal statutes organized topically, and Title 18 is the main place to find federal criminal laws). In particular, it is codified as 18 U.S.C. Section 1030.

For some of you, this will be your first time to really examine a criminal statute. You may be surprised by how convoluted it seems to be once you begin reading it closely. I assure you: You can handle it, so long as you take your time parsing the language.

The part we want to focus on here is the first subsection, 1030(a). As you will see, it actually contains seven *separate* criminal offenses. I reprint them below, but you can also find the full text of the statute (including the relevant definitions) here. Read them through one time, slowly, and then proceed to the questions below.

Section 1030(a) imposes felony liability on whomever:

> **(1)** having knowingly accessed a computer without authorization or exceeding authorized access, and by means of such conduct having obtained information that has been determined by the United States Government pursuant to an Executive order or statute to require protection against unauthorized disclosure for reasons of national defense or foreign relations, or any restricted data, as defined in paragraph y. of section 11 of the Atomic Energy Act of 1954, with reason to believe that such information so obtained could be used to the injury of the United States, or to the advantage of any foreign nation willfully communicates, delivers, transmits, or causes to be communicated, delivered, or transmitted, or attempts to communicate, deliver, transmit or cause to be communicated, delivered, or transmitted the same to any person not entitled to receive it, or willfully retains the same and fails to deliver it to the officer or employee of the United States entitled to receive it;

(2) intentionally accesses a computer without authorization or exceeds authorized access, and thereby obtains—

> **(A)** information contained in a financial record of a financial institution, or of a card issuer as defined in section 1602(n) of title 15, or contained in a file of a consumer reporting agency on a consumer, as such terms are defined in the Fair Credit Reporting Act (15 U.S.C. 1681 et seq.);
>
> **(B)** information from any department or agency of the United States; or
>
> **(C)** information from any protected computer;

(3) intentionally, without authorization to access any nonpublic computer of a department or agency of the United States, accesses such a computer of that department or agency that is exclusively for the use of the Government of the United States or, in the case of a computer not exclusively for such use, is used by or for the Government of the United States and such conduct affects that use by or for the Government of the United States;

(4) knowingly and with intent to defraud, accesses a protected computer without authorization, or exceeds authorized access, and by means of such conduct furthers the intended fraud and obtains anything of value, unless the object of the fraud and the thing obtained consists only of the use of the computer and the value of such use is not more than $5,000 in any 1-year period;

(5)

> **(A)** knowingly causes the transmission of a program, information, code, or command, and as a result of such conduct, intentionally causes damage without authorization, to a protected computer;
>
> **(B)** intentionally accesses a protected computer without authorization, and as a result of such conduct, recklessly causes damage; or
>
> **(C)** intentionally accesses a protected computer without authorization, and as a result of such conduct, causes damage and loss.

(6) knowingly and with intent to defraud traffics (as defined in section 1029) in any password or similar information through which a computer may be accessed without authorization, if—

> **(A)** such trafficking affects interstate or foreign commerce; or
>
> **(B)** such computer is used by or for the Government of the United States;

(7) with intent to extort from any person any money or other thing of value, transmits in interstate or foreign commerce any communication containing any—

> **(A)** threat to cause damage to a protected computer;
>
> **(B)** threat to obtain information from a protected computer without authorization or in excess of authorization or to impair the confidentiality of information obtained from a protected computer without authorization or by exceeding authorized access; or
>
> **(C)** demand or request for money or other thing of value in relation to damage to a protected computer, where such damage was caused to facilitate the extortion. . . .

Also note that the statute contains several definitions in 1030(e), too many to list here. Here are the highlights:

> **(1)** the term "computer" means an electronic, magnetic, optical, electrochemical, or other high speed data processing device performing logical, arithmetic, or storage functions, and includes any data storage facility or communications facility directly related to or operating in conjunction with such device, but such term does not include an automated typewriter or typesetter, a portable hand held calculator, or other similar device;
>
> **(2)** the term "protected computer" means a computer—
>
> > **(A)** exclusively for the use of a financial institution or the United States Government, or, in the case of a computer not exclusively for such use, used by or for a financial institution or the United States Government and the conduct constituting the offense affects that use by or for the financial institution or the Government; or
> >
> > **(B)** which is used in or affecting interstate or foreign commerce or communication, including a computer located outside the United States that is used in a manner that affects interstate or foreign commerce or communication of the United States;
>
> **(6)** the term "exceeds authorized access" means to access a computer with authorization and to use such access to obtain or alter information in the computer that the accesser is not entitled so to obtain or alter;
>
> **(8)** the term "damage" means any impairment to the integrity or availability of data, a program, a system, or information;
>
> **(11)** the term "loss" means any reasonable cost to any victim, including the cost of responding to an offense, conducting a damage assessment, and restoring the data, program, system, or information to its condition prior to the offense, and any revenue lost, cost incurred, or other consequential damages incurred because of interruption of service; and
>
> **(12)** the term "person" means any individual, firm, corporation, educational institution, financial institution, governmental entity, or legal or other entity.

To help you make sense of the statute, consider the following:

- Make a list, chart, flash cards, or something else that helps you break down the key elements for each of these separate crimes. Start by giving each one a label, in your own words, that fairly describes the unique aspect of that particular provision. Then list the elements necessary for an offense, including the necessary actions, the required mental state, and any other conditions mentioned.
- Do any of the provisions seem problematic to you? Be prepared to explain why.
- Which provisions seem must relevant to a classic "hacking" scenario in which someone uses an exploit to gain unauthorized access to someone else's computer, and thereafter exfiltrates data from it?
- Is the CFAA *only* concerned with such hacking scenarios?

Notice that the CFAA goes on, in subsection 1030(b), to criminalize conspiracies and attempts to commit the offenses listed in 1030(a).

3. CFAA Case Studies

As we read through various cases where individuals were charged with violating the CFAA, bear in mind that the statute has been amended several times throughout the years. For some of these cases, the charged offenses will differ slightly from the modern ones. As you read, consider whether it would be easier or harder to charge the same individual with a CFAA violation today, and whether any other offenses might apply to the same conduct.

The first big CFAA prosecution involved the ground-breaking—and largely accidental—"Morris Worm."

Read about the underlying events here, and then read the court opinion affirming his conviction—*United States v. Morris*, 928 F.2d 504 (2d Cir. 1991).

- Do you agree that Morris violated the CFAA—either today or as it was written then?
- What were the best legal arguments for and against prosecuting Morris?
- Does the prosecution make normative sense to you? In other words, does it make sense to prosecute people like Morris?

The controversy surrounding the Morris case was nothing compared to that generated by the later prosecution of Aaron Swartz.

Read this article about the Swartz prosecution, and consider the same questions from the Morris case.

- Do you agree that Swartz violated the CFAA—either today or as it was written then?
- What were the best legal arguments for and against prosecuting Swartz?
- Does the prosecution make normative sense to you? In other words, does it make sense to prosecute people like Swartz?

Another much-discussed example concerned David Nosal, who once worked for the executive search-and-recruitment firm Korn/Ferry and then left to start a competitor. It was not exactly a clean separation, however, from an information security perspective.

Read about the CFAA charges in Nosal's case, the issues they raised, and the outcome as explained by the "*en banc*" Ninth Circuit Court of Appeals in *United States v. Nosal*, 676 F.3d 854 (2012).

- What was the government's theory of how Nosal violated the CFAA? Which subsections did the government argue he violated?
- What was Nosal's counterargument?
- What did the court decide, and what was its reasoning? Do you agree?

The government on remand from that decision re-tried Nosal, this time advancing a modified theory as to why his conduct violated the CFAA.

Read this article on what happened next and consider the following questions:

- What was the government's revised theory?
- Did the court accept it this time?
- Are you persuaded by the government's new argument?

4. Other Criminal Statutes

A. Other Relevant Criminal Laws

The CFAA isn't the only tool in the toolbox for federal prosecutors dealing with cybercrime. There are some statutes of more-general applicability that often fit well with hacking scenarios, and there also are some highly-tailored statutes to consider.

The most-relevant of the generally-applicable criminal laws, in this respect, is the "wire fraud" statute—18 U.S.C. 1343:

> Whoever, having devised or intending to devise any scheme or artifice to defraud, or for obtaining money or property by means of false or fraudulent pretenses, representations, or promises, transmits or causes to be transmitted by means of wire, radio, or television communication in interstate or foreign commerce, any writings, signs, signals, pictures, or sounds for the purpose of executing such scheme or artifice, shall be fined under this title or imprisoned not more than 20 years, or both. If the violation occurs in relation to, or involving any benefit authorized, transported, transmitted, transferred, disbursed, or paid in connection with, a presidentially declared major disaster or emergency (as those terms are defined in section 102 of the Robert T. Stafford Disaster Relief and Emergency Assistance Act (42 U.S.C. 5122)), or affects a financial institution, such person shall be fined not more than $1,000,000 or imprisoned not more than 30 years, or both.

Consider the following questions:

- How does the wire fraud statute differ from the CFAA?
- Why might a prosecutor find this statute handy?

To get a further sense of what a wire fraud prosecution linked to hacking might look like, check out this article about an unusual wire fraud prosecution. Gooooooooooooooal!

Apart from "wire fraud," there are several other, more-specific fraud statutes. While I do not intend for you to learn the particulars with them (as you will with the CFAA and wire fraud), I do want you to be familiar with the general idea behind them. So skim the following information on identify fraud, identity theft, and access device fraud sufficiently to be able to articulate what they forbid.

Identity fraud, 18 USC 1028, allows the government to prosecute an individual who:

> **(1)** knowingly and without lawful authority produces an identification document, authentication feature, or a false identification document;
>
> **(2)** knowingly transfers an identification document, authentication feature, or a false identification document knowing that such document or feature was stolen or produced without lawful authority;
>
> **(3)** knowingly possesses with intent to use unlawfully or transfer unlawfully five or more identification documents (other than those issued lawfully for the use of the possessor), authentication features, or false identification documents;
>
> **(4)** knowingly possesses an identification document (other than one issued lawfully for the use of the possessor), authentication feature, or a false identification document, with the intent such document or feature be used to defraud the United States;
>
> **(5)** knowingly produces, transfers, or possesses a document-making implement or authentication feature with the intent such document-making implement or authentication feature will be used in the production of a false identification document or another document-making implement or authentication feature which will be so used;
>
> **(6)** knowingly possesses an identification document or authentication feature that is or appears to be an identification document or authentication feature of the United States or a sponsoring entity of an event designated as a special event of national significance which is stolen or produced without lawful authority knowing that such document or feature was stolen or produced without such authority;
>
> **(7)** knowingly transfers, possesses, or uses, without lawful authority, a means of identification of another person with the intent to commit, or to aid or abet, or in connection with, any unlawful activity that constitutes a violation of Federal law, or that constitutes a felony under any applicable State or local law; or

> **(8)** knowingly traffics in false or actual authentication features for use in false identification documents, document-making implements, or means of identification;

Only when:

> **(1)** the identification document, authentication feature, or false identification document is or appears to be issued by or under the authority of the United States or a sponsoring entity of an event designated as a special event of national significance or the document-making implement is designed or suited for making such an identification document, authentication feature, or false identification document;
>
> **(2)** the offense is an offense under subsection (a)(4) of this section; or
>
> **(3)** either—
>
> > **(A)** the production, transfer, possession, or use prohibited by this section is in or affects interstate or foreign commerce, including the transfer of a document by electronic means; or
> >
> > **(B)** the means of identification, identification document, false identification document, or document-making implement is transported in the mail in the course of the production, transfer, possession, or use prohibited by this section.

Identity theft, 18 USC 1028A, allows the government to prosecute an individual who:

> during and in relation to any felony violation enumerated in subsection (c), knowingly transfers, possesses, or uses, without lawful authority, a means of identification of another person . . .

Subsection (c) lists several offenses contained in other statutes, including theft and embezzlement; "false personification of citizenship"; "false statements in connection with the acquisition of a firearm"; fraud; mail, bank, and wire fraud; several immigration-related offenses; "obtaining customer information by false pretenses"; and "false statements relating to" Social Security programs.

Access device fraud, 18 USC 1029, allows the government to prosecute an individual who, in affecting interstate commerce:

> **(1)** knowingly and with intent to defraud produces, uses, or traffics in one or more counterfeit access devices;
>
> **(2)** knowingly and with intent to defraud traffics in or uses one or more unauthorized access devices during any one-year period, and by such conduct obtains anything of value aggregating $1,000 or more during that period;
>
> **(3)** knowingly and with intent to defraud possesses fifteen or more devices which are counterfeit or unauthorized access devices;
>
> **(4)** knowingly, and with intent to defraud, produces, traffics in, has control or custody of, or possesses device-making equipment;

> **(5)** knowingly and with intent to defraud effects transactions, with 1 or more access devices issued to another person or persons, to receive payment or any other thing of value during any 1-year period the aggregate value of which is equal to or greater than $1,000;
>
> **(6)** without the authorization of the issuer of the access device, knowingly and with intent to defraud solicits a person for the purpose of—
>
> **(A)** offering an access device; or
>
> **(B)** selling information regarding or an application to obtain an access device;
>
> **(7)** knowingly and with intent to defraud uses, produces, traffics in, has control or custody of, or possesses a telecommunications instrument that has been modified or altered to obtain unauthorized use of telecommunications services;
>
> **(8)** knowingly and with intent to defraud uses, produces, traffics in, has control or custody of, or possesses a scanning receiver;
>
> **(9)** knowingly uses, produces, traffics in, has control or custody of, or possesses hardware or software, knowing it has been configured to insert or modify telecommunication identifying information associated with or contained in a telecommunications instrument so that such instrument may be used to obtain telecommunications service without authorization; or
>
> **(10)** without the authorization of the credit card system member or its agent, knowingly and with intent to defraud causes or arranges for another person to present to the member or its agent, for payment, 1 or more evidences or records of transactions made by an access device;

The statute also defines an "access device" as:

> any card, plate, code, account number, electronic serial number, mobile identification number, personal identification number, or other telecommunications service, equipment, or instrument identifier, or other means of account access that can be used, alone or in conjunction with another access device, to obtain money, goods, services, or any other thing of value, or that can be used to initiate a transfer of funds (other than a transfer originated solely by paper instrument)

For a fascinating case showing how a variety of these statutes might be used in combination, we will discuss the prosecution of Roman Zeleznev (aka "Track2").

Read about it here and consider the following questions:

- This case turned out well for the government; do you think it indicates that similar success is possible in most such cases? Read this and this for a glimpse of some unusual complications.
- What made this case harder than normal? What factors explain how DOJ prevailed anyway? Does this show DOJ can generally prevail in similar cases?

Note: There are other criminal laws that are important in this space, but that we will not explore in the interest of moving along to other topics. For the record, however, I'll still proceed to name a few of them: 18 USC 641 (theft of government property); 18 USC 2511 (unauthorized interception of communications); 18 USC 2701 (unauthorized accessing of stored communications); and 18 USC 793–798 (various provisions relating to espionage and protection of defense information). There also is 17 USC 1201–1205, aka the Digital Millennium Copyright Act ("DMCA").

B. State Criminal Laws

States have statutes analogous to the CFAA. For an overview of the relevant Texas statute, including observations on how it differs from CFAA in certain respects, read this article. For a sense of the state agency responsible for computer crime investigations, by the way, read here.

> Be prepared to articulate whether / how the Texas statute differs from the CFAA.

C. International Cybercrime Enforcement

The United States is party to the "Budapest Convention on Cybercrime"—an international treaty promoting cooperation between nations in combating cybercrime.

> Skim its provisions to get a rough sense of what it is trying to accomplish.
> - What do the parties to this treaty actually promise to do that seems genuinely significant?
> - Why do you suppose Russia, China, and Iran are not parties to this treaty?

5. Civil Liability under the CFAA

It turns out the CFAA is not just a *criminal* statute, but also a *civil liability* statute—that is, it also creates a "private right of action" enabling suits for money damages in certain circumstances.

Read 10 USC 1030(g). For such a short subsection, there is a LOT going on here!

- First, there is a complicated precondition for exercising that right to sue, to the effect that only certain conduct counts.
    - Can you unpack the statutory cross-references and explain, in plain English, when someone is allowed to sue?
    - Can you explain what this leaves out, and why Congress might have gone to such trouble to draw this line?
    - Do you agree with this approach?
- Next, there's a sentence that limits the plaintiff to "economic damages" for a certain type of case.
    - What does this mean, what type of case counts, and what explains all this?
- We'll skip over the statute of limitations (that is, the part that says you only get two years to sue). That brings us to the last sentence of 1030(g).
    - What precisely does this last sentence do?
    - What are the arguments for and against this provision.

Here's the basic provision:

> **(g)** Any person who suffers damage or loss by reason of a violation of this section may maintain a civil action against the violator to obtain compensatory damages and injunctive relief or other equitable relief. A civil action for a violation of this section may be brought only if the conduct involves 1 of the factors set forth in subclauses (I), (II), (III), (IV), or (V) of subsection (c)(4)(A)(i). Damages for a violation involving only conduct described in subsection (c)(4)(A)(i)(I) are limited to economic damages. No action may be brought under this subsection unless such action is begun within 2 years of the date of the act complained of or the date of the discovery of the damage. No action may be brought under this subsection for the negligent design or manufacture of computer hardware, computer software, or firmware.

The factors set out in subsection (c)(4)(A)(i) refer to an offense (or attempt) causing (or attempting to cause):

> **(I)** loss to 1 or more persons during any 1-year period (and, for purposes of an investigation, prosecution, or other proceeding brought by the United States only, loss resulting from a related course of conduct affecting 1 or more other protected computers) aggregating at least $5,000 in value;
>
> **(II)** the modification or impairment, or potential modification or impairment, of the medical examination, diagnosis, treatment, or care of 1 or more individuals;
>
> **(III)** physical injury to any person;
>
> **(IV)** a threat to public health or safety;
>
> **(V)** damage affecting a computer used by or for an entity of the United States Government in furtherance of the administration of justice, national defense, or national security; or

> **(VI)** damage affecting 10 or more protected computers during any 1-year period;

Note: Factor (VI) is not covered by section (g)—in other words, damage affecting 10 or more protected computers in one year will not alone give rise to civil liability.

Social media companies like Facebook and LinkedIn at times turn to the CFAA's civil liability provision in an effort to stop other companies from collecting information from public-facing parts of their sites. Whether and when such conduct violates the CFAA is the subject of much debate.

Read about the suit Facebook filed against "Power Ventures" here and here.

- How did Power Ventures make use of data found on Facebook?
- What were Facebook's arguments about how this violated the CFAA? How did Power Ventures respond?
- How did the courts resolve the matter? Do you agree?

Note: The Supreme Court recently refused to review the appellate court's ruling, ending the case.

HiQ v. LinkedIn is a similar, recent case.

Read the following article about the litigation, and then read the Ninth Circuit's recent opinion in the case.

- How did HiQ make use of data found on LinkedIn?
- What were LinkedIn's arguments about how this violated the CFAA? How did HiQ respond?
- How did the Ninth Circuit resolve the matter? Do you agree?
- What are the larger policy stakes at play in this case?

6. What if the Attacker is a Foreign Government? (I)

A. Key Concepts

Up until now, we have proceeded from the assumption that instances of unauthorized access involved run-of-the-mill criminal or tortious activity conducted by private individuals or organizations. But sometimes the perpetrator is acting on behalf of a foreign government. Governments hack for many different (and sometimes overlapping) reasons, and we need to introduce these possibilities before turning to the questions that arise when we consider how the U.S. government attempts to impose costs on foreign governments in these situations.

The primary purposes behind government hacking, in no particular order, include:

- **Law Enforcement:** A government may engage in hacking to advance its own law enforcement interests. This could involve hacking to investigate or gather evidence on a crime that has already been committed, or perhaps executing a kind of sting operation.

- **Crime:** Some regimes are desperate for cash. Private persons are not the only ones who might hack for financial gain.

- **Information Collection:** Most states are in the business of stealing secrets in order to inform decision-making or to advance other goals, though states vary widely in their capacity to actually do this effectively. It is an ancient art, one that always has involved both technical and non-technical means. As more and more information and communications have gone digital, hacking has become ever more central to intelligence collection. Often we call this "spying" or "espionage," terms that call to mind images of civilian agencies stealing secrets for the benefit of a government (or, in the practice of some states—though *not* the United States—for the benefit of state-controlled or state-favored private enterprises). But civilian agencies are not the only ones that engage in surreptitious information collection. When conducted by the military, we sometimes refer to this activity as intelligence, surveillance, and reconnaissance ("ISR"). ISR has connotations of informing tactical, operational, or even strategic military planning. Whatever the label, though, the bottom line is that hacking is an increasingly-necessary aspect of stealing secrets.

- **Covert Action:** As our review of the CFAA underscored, the general label of "hacking" encompasses more than just unauthorized access to steal information; sometimes the access is sought in order to alter or destroy data or to cause harm to a system controlled or impacted by that data. When a government pursues that approach, a question arises regarding how to categorize the activity. Sometimes such activity will be part of an armed conflict, and we will say more about that in just a moment. For now, what matters is that not every such hack occurs in the context of armed conflict; indeed, most do not. And yet they are not instances of espionage, either. So what are they? Well, if the government involved is trying to keep its role secret, then the best answer usually will be "covert action." There are some complicated nuances here, at least within the U.S. legal system, when the government entity involved is a military entity—but we will save that for later in the course. Covert action can encompass a wide-range of cyber operations, from information operations (propaganda, disinformation, etc.) to efforts to create damaging physical effects (sabotage).

- **Armed Conflict:** Though journalists and others routinely refer to "cyber attacks" and "cyberwar" when talking about hostile foreign cyber activities targeting U.S. systems, the fact is that these actions rarely actually concern genuine armed conflict involving the United States. But they certainly *could*, and sometimes they really do. And so the threshold question you must ask is: Is there already a relevant state of armed conflict, or could this action on its own engender one? If not, then it is better not to talk in terms of war and combat; covert action may be the better label.

- **Preparation of the Battlefield:** This is military jargon for the idea that it is at times useful or even necessary for the armed forces to take certain actions in advance of potential hostilities—sometimes *far* in advance—in order to be in a better position to carry out certain operations later (that is, if and when an armed conflict actually begins). In the physical world, for example, special operators might enter enemy territory prior to a conflict in order to determine optimal routes, preposition supplies, and so forth. So too, then, with cyber operations. In order to be able to take an action involving a targeted system later, it may be wise or even necessary to establish access to that system now. Of course, in that case, it's best to remain undetected, lest that "preparation of the

battlefield" go to waste. But what if one actually wants to be detected? That leads us to the distinct concept of a "hold-at-risk" strategy.

- **Hold-at-Risk:** This perhaps unfamiliar phrase is a shorthand for a simple idea. It refers to the idea that one might want to demonstrate to a rival or prospective opponent—in a very credible way—that one has the capacity to cause damage to something that they value. That is, the idea is to prove that you are holding something they value "at risk," and that the other side had best not forget this when interacting with you in other settings. Simply put, a "hold-at-risk" strategy is an effort to improve your deterrence posture in relation to an opponent, thus impacting their calculations and actions in a way that is favorable to you. In this sense, it is akin to a "show of force" in which a government puts equipment or personnel, quite visibly, in geographic position to carry out certain operations (e.g., positioning an aircraft carrier nearby). In the cyber context, penetrating a system that contains valuable data or controls a valuable system—and allowing the other side to detect that you have done this—is a way of signaling that you truly do have the means to harm that data or system (and perhaps others as well).

- **The Indeterminacy and Multiplicity Problems:** Here's the most important point of them all: In many instances, it is not easy for a defender to tell which of these aims might explain why someone has hacked into a particular system. To be sure, it can become clear enough once the hacker begins making use of that access in order to do certain things. But because all of the aforementioned motivations for state-sponsored hacking begin with social engineering or malware in order to gain unauthorized access to a system, a defender who has detected such an intrusion may be left with little basis for predicting what the intruder intends. Sometimes the context will help, of course, and eventually time will tell. In rare instances, moreover, external sources of information might shed useful light too. But in the meantime, the defender is left to make the best guess possible in the circumstances. Note, too, that the intruder might have in mind one purpose at one time, and may switch to a different purpose later.

Bearing the above in mind, perform the following exercise:

- Pick a foreign power with whom the United States has particularly bad relations. Imagine you are in a position of authority and trust in that government, and imagine a concrete example of an American target that you might like to have your government penetrate for each of the purposes mentioned above.
- Be able to explain what your country gets from each scenario, as well as the offsetting risks that might give you pause.

7. What if the Attacker is a Foreign Government? (II)

A. A Framework for Thinking about Threat Reduction

Conversations about the threat foreign governments may pose to networks in the United States (including not just threats associated with formal parts of those governments, like their militaries and intelligence services, but also private persons/organizations that may act on behalf of those governments) often are framed in terms of "deterrence," "escalation risk," and other familiar concepts from the international relations and security literatures. And rightly so. Before exploring

those concepts in detail, however, it might help to spend a moment considering, at a high level of generality, how such concepts fit into a larger picture.

Let's say you are the President of the United States, and you and your advisors are formulating a strategy in response to your belief that a foreign government—let's say Iran—might take an action you view as a serious threat to U.S. interests. That action could involve the use of an existing capability in some unwelcome way—a use of military force, an intervention in the oil market, etc. Or it might involve an attempt to acquire some new capability that would make the foreign state a greater threat in the future—a nuclear bomb, for example. The point is, your overarching goal is to minimize the net danger.

To achieve that overarching goal, there are at least three subsidiary strategies you might consider (they are not mutually exclusive):

Disruption	Deterrence	Defense
Prevent the other state from becoming capable of taking the undesirable action. Or, if it has capability already, destroy (or at least degrade) it.	Convince them not to take the undesirable action by maximizing their expected costs and minimizing their expected benefits.	Minimize the harm you would suffer if the undesirable action occurs by maximizing relevant defenses and establishing resiliency.

Note that the disruption, deterrence, and defense strategies can relate to one another. For example, if you build strong defenses and your adversary knows this, this may cause an increase in their expected costs of the action (for they may conclude they must put more of their own resources into the effort) or a decrease in their expected benefits (for they may have to revise their odds of success). Either way, their cost-benefit assessment will be less appealing, and the degree of deterrent persuasion increased.

B. Notes on the U.S. Defensive Posture in Cyberspace: America the Vulnerable?

Read this article about how extensive "digitalization" of the United States results in strategic vulnerabilities that, in turn, cast a shadow over U.S. policymaking and operational decision-making in response to adversary actions in cyberspace.

As you read, consider the following:

- What are the six vulnerabilities described, and precisely why it should matter to U.S. policymakers pondering their options in response to hostile foreign cyber activity?
- What larger lessons do you draw?

C. Key Terms Relating to Deterrence

The following set of readings are designed to introduce you to some of the terminology related to deterrence.

Read this article and be prepared to define these terms:

- Deterrence
- Cross-domain deterrence
- Within-domain deterrence
- Escalation
- Escalation risk
- Escalation dominance

Next, read this article.

Consider the following questions:

- Should the government always make public that it has taken an action in response to another state's hacking? Can deterrence work without public claims of that kind?
- In answering those questions, give thought to the different possible "audiences" for such actions. Obviously, one would be the foreign government to which the U.S. is responding. But who else might be watching?

On "attribution" in the cyberspace context, read this and this.

Consider these questions:

- Can you define "attribution" in this setting?
- Why do some claim attribution is especially difficult in the cyberspace context, and why would that be different than, say, nuclear weapons?
- What impact does such difficulty mean as to decision-making in particular cases?

8. What if the Attacker is a Foreign Government (III)

There are many ways in which governments can impose costs on foreign states following a hacking—including prosecution and economic sanctions. For this class, we'll look at when it makes sense to use some or several of these tools.

Prosecution

Consider the following account from Associate Deputy Attorney General Sujit Raman from a speech delivered in May 2019:

> For many years, we have targeted and successfully disrupted transnational criminal syndicates engaged in cybercrime, as well as the digital infrastructure those actors employ. More recently, we began publicly charging foreign state actors whose malicious cyber activity broke U.S. law. It is fair to ask why we devote significant resources to prosecuting state actors whom we may never bring to the United States to face justice. And it is fair to ask why we shifted from an approach that relied mainly on intelligence collection and diplomacy to one that includes a law enforcement response. As I will explain, prosecutions of state-sponsored malicious cyber activity serve

an important purpose, even if we cannot guarantee that we will be able to produce in court every individual involved.

Since the indictment of Chinese PLA officers in 2014, the Department of Justice has remained focused on state-sponsored criminal activity that targets U.S. companies. We are also now focused on activity that targets the U.S. political process. In the past two years, the Department brought more national security cyber cases against criminals acting on behalf of our major adversaries than in the previous five years.

There are several reasons for the increasing prosecutions. The main one is that we are following the threat, just as we did in responding to the threat of terrorism. As I've explained, nation states are engaged in activity that victimizes individuals and companies in the United States, violates U.S. law, and departs from international norms of responsible state behavior—norms that benefit all nations. Our criminal cases reflect our adversaries' efforts to harm our companies and our nation.

Second, the increasing number of national security cyber cases reinforces the lesson that our adversaries' conduct lies outside the norms of responsible behavior. The actions we highlight in indictments are not legitimate state-craft. They are crimes without justification in international relations. I will say more about that in a moment.

Third, our cases reflect our increasingly sophisticated ability to attribute this criminal conduct to the individuals and states involved. This ability is closely related to my second point, because it shows the commitment of our law enforcement and intelligence agencies to work closely together, while protecting intelligence sources and methods. These partnerships, which were forged in the counterterrorism context, serve to solidify the consensus that a law enforcement response to malign nation-state cyber activity makes sense.

In bringing these cases, we are guided by six basic principles.

First, the Department has a duty to enforce our laws and protect our people. We cannot refuse to act when foreign state actors violate our criminal laws, transgress established norms, and victimize our citizens. That is especially true when such crimes often represent severe violations of the victim's privacy rights, and can have lasting, damaging impact. The Department has an obligation to work toward a future where our citizens can live and conduct their business with confidence in the integrity of their information and institutions.

Second, attribution is the key to deterrence. "Without attribution, there will be no consequences . . . and thus, no deterrence."[14] Attribution through the criminal justice system escalates the stakes for state-sponsored activity in a way that a press release or a public statement alone would not. We have on occasion obtained custody of foreign criminal defendants. Our indictments limit their travel. And the prospect of criminal indictment can help deter some cyber actors from engaging in such conduct in the first place. This can make it more difficult for states to recruit the manpower and resources for cyber-attacks, and raise the cost of engaging in malicious cyber activity.

Third, attribution through the criminal justice system is a powerful way to expose state conduct that violates norms of responsible behavior. It complicates our adversaries' attempts to feign ignorance of illegal acts they thought could be taken in secret, or to hide behind public denials. Our cases are governed by well-known policies relating to the conduct of all federal prosecutors. An indictment is brought by a grand jury, under established procedures; charges are brought only when the facts and law justify it. The

> allegations in our indictments are thorough and detailed, and we can prove them in a courtroom, using admissible evidence, at proof beyond any reasonable doubt. For all these reasons, criminal indictments are among the most powerful statements we can make as a government.
>
> Fourth, unsealed indictments promote transparency. There will always be cases in which our ability to expose malicious cyber activity is limited by our obligation to protect intelligence sources and methods or sensitive ongoing investigations. But where we are able to do so, we strive to expose such schemes to the American people and the international community. Attribution through detailed indictments educates the public about our adversaries' efforts and methods to spread disinformation, steal commercial technology, and target computer networks.
>
> Fifth, although our goal is to hold accountable in court those we charge with trade theft or cyber crimes, our investigations can provide critical support for the use of civil, diplomatic, economic, and military tools. Some thoughtful critics have criticized the Department's so-called "name and shame" strategy on the theory that our indictments have failed to stanch the activity. But you can't separate our indictments from the broader array of tools our government now uses to counter malign cyber activity. These include freezing assets or prohibiting transactions, or adding companies to the Department of Commerce Entity List. As the National Security Advisor has confirmed, it also includes "undertaking offensive cyber operations"[15] aimed at defending our national interests. Our tools also include other authorities that can block criminals' assets, restrict their access to the banking system, and prohibit them from freely engaging in trade. We developed this approach to address terrorism and terrorist financing. We are applying it to the cyber threat, as well.
>
> Finally, by using public law to emphasize the need to protect private U.S. victims against nation-state actors, we help develop the framework of public-private cooperation that is critical to cybersecurity. The Department tries to show through our actions how we can help companies respond to nation-state threats they cannot face alone, in a way that respects their status as victims. The Department has developed strong relationships with the private sector based on our aggressive pursuit of criminal nation-state conduct, ranging from cyber theft to information operations using third-party social media platforms.
>
> No one seriously suggests that we can prosecute our way out of this problem. But to dismiss the role that federal law enforcement plays in the government's shared fight against cyber-enabled threats is to unfairly discount—and diminish—our nation's powerful commitment to the rule of law, both within our borders and without.

Economic Sanctions

There are many other tools to bear in mind, including economic sanctions. This is a vast and important topic, the breadth of which is well beyond the scope of our course. Nonetheless, this section conveys what you should understand at a minimum.

When we talk of "sanctions" in this setting, we are referring to the ability of the U.S. government either to freeze the U.S.-based assets of a foreign person, organization, or government, or to declare some or all transactions with the sanctioned party unlawful—meaning no purchases, sales, trade, donations, services, exchanges, etc.

Congress at times has passed laws that directly impose sanctions, but it is much more common these days for Congress to delegate to the President the authority to impose sanctions based on certain criteria Congress sets. For example, Congress recently enacted the Countering America's Adversaries Through Sanctions Act (CAATSA, pronounced "Cats-uh" or "Cots-uh"), which among other things calls for sanctions in response to interference with U.S. elections. And definitely be aware of the International Emergency Economic Powers Act (IEEPA, pronounced "Eye-EEP-uh"), which since the 1970s has served as a broad delegation of authority for the president to sanction foreign entities so long as the president has publicly declared the existence of a "national emergency" relating to a foreign affairs matter and the sanctions are related to that situation.

Note: Don't be misled by the seeming gravity of declaring a national emergency; the public over time has proven to be largely uninterested when such declarations occur, and thus it has proven relatively easy to declare them when deemed useful.

By and large, presidential authority to issue sanctions under these and similar statutes ends up being delegated, through an executive order, to the Treasury Department. The body within Treasury that manages sanctions is the Office of Foreign Assets Control ("OFAC"). Periodically, OFAC will announce new entities or individuals to be sanctioned under one or more of the currently active sanction regimes.

What makes people comply with sanctions? Criminal penalties for violating them, derived from the statutes that created the sanction rules in the first place. Can you sanction someone for their violation of other sanctions? Yes, that's called "secondary sanctions." This is a hot topic vis-à-vis foreign companies that want to do business with foreign entities (such as certain Iranian entities) that are themselves the subject of sanctions.

Unilateral sanctions (that is, those imposed only by the United States) can have an impact, but multilateral sanctions of course can have a greater impact. To get other states to follow suit, the U.S. government can try diplomatic persuasion. To actually compel other governments to follow suit? That requires a U.N. Security Council Resolution, which is no easy thing to obtain given the constellation of competing national interests the Council represents (and the fact that China, Russia, the UK, France, and the United States all have permanent authority to veto UNSC action).

Can you think of other tools the U.S. government can bring to bear to impose costs? What about "carrots" the U.S. government can offer instead? Generate a list, and for each tool, answer the following questions:

- To what extent would the foreign government view the use of that tool as undesirable to it?
- To what extent is it possible for the U.S. government to actually bring that tool to bear (and would the other government likely understand this)?
- To what extent does the relevant U.S. decision-maker have the will to use that tool? And, more to the point, is the other government likely to know that?
- While our discussion of these tools has focused on deterrence, would any of the tools on your list also be useful for disruption?

9. What if the Attacker is a Foreign Government? (IV)

Our aim at the outset of this class is to use a handful of case studies in order to understand deterrence dynamics in relation to cyberspace.

A. Russia

Read this New York Times account from December 2016, which focuses on Russian election interference in 2016, and this one on possible responses. Next, have a look at this indictment of various Russians officers issued by a grand jury in July 2018.

As you read, consider the following questions:

- How did hacking in this context relate to a larger "information operation"? What lessons does this episode suggest about vulnerability to spear-phishing?
- How do you assess the response of (i) the FBI in particular and (ii) the U.S. government more generally?
- What insights did you gather regarding the entities that conduct such activities for the Russian government? How would you characterize what Russia did here? (Crime? Espionage? Covert action? Some combination, or something else?)
- What would you have done differently had you been president?
- And, finally, is any of this actually beyond the pale, in the sense that you would not want to see the United States doing the same thing (and do you see how that is a different question from asking whether the United States should do what it can to stop such actions from succeeding against it)?

Read this Washington Post story regarding another Russia incident, the resulting indictment, and the latest developments in the case.

How do you assess the effectiveness of the U.S. government response in this instance? Can the same model reliably be applied elsewhere?

B. China

Read this Christian Science Monitor piece examining Chinese-government sponsored hacking against U.S. targets (public and private) in recent years.

Consider the following questions:

- How does Chinese-government sponsored hacking differ from the Russian activities descried above, and what follows from this as a matter of policy?
- Should it be off-limits to hack businesses in hopes of providing competitive advantages for your own nation's companies (and does it really matter if the companies in question, on either end, are formally owned in part or in whole by the state)?
- Should the United States have done more to respond to these hacks?

The Obama administration surprised many observers when it brought criminal charges against a group of PLA hackers (*United States v. Dong* (W.D. Pa.). Read the indictment, as well as this story and this story about the case.

Analysis of the impact of this effort has been conflicted. Compare this account and this account.

- What lessons if any do you draw from this? Is prosecution an effective approach? Scalable?
- Does news of this recent arrest –possibly linked to the famous OPM hack—change your view?

Prosecution is not the only tool available, of course. Read this Executive Order from President Obama, which in April 2015 established a system for sanctioning the beneficiaries of cyber-espionage used for commercial advantage. Read more here and here, too.

What are the pros and cons of this approach?

Eventually, the U.S. and Chinese government struck a deal, of sorts.

Read this recent account, and consider the following questions:

- What was the deal, and has it helped?
- What lessons do you draw?

B. Encouraging Potential Victims to Defend Better

Minimizing unauthorized (and excessive) access is not just a function of imposing painful consequences on intruders. It is also a function of making it harder for them to succeed when they *do* make an attempt—*i.e.*, it is a function of improving defense, too.

Recall how we distinguished disruption, deterrence, and defense in the readings above. Improving defense will, at the margins, prevent some intrusions in the direct sense that some attacks that might otherwise have succeeded will now fail. Moreover, even for attackers who are able and willing to overcome the improved defense, the improvement increases the attacker's costs, thereby making the effort marginally less attractive (by reducing prospective return on investment) and perhaps even causing the attacker to reduce the scope of their activities due to resulting resource constraints. Sometimes this is called "deterrence by denial." Incentivizing potential victims (whether they are private or public entities or individuals) to improve their defenses on a systemic basis thus can serve important goals of cybersecurity policy.

Of course, most potential victims already have at least *some* incentive to develop and improve defenses, even absent any form of external intervention to encourage them further. Some have trade secrets to protect. Most desire to keep at least some things private. Some need to keep customers happy. And so forth. As a result, we can safely assume there will be *some* defensive activity even if *no* external forces intervene to encourage such steps. It's rather like the situation

of a building owner. Most building owners would take at least *some* steps to make the building safe even if there were no building codes, insurers, or plaintiff's lawyers with which to contend. We might call this the "natural level" of investment in defense.

Is the natural level good enough? In the building-safety context, society has answered that question with a resounding no. Governments, insurers, and litigators intervene in all sorts of ways to spur further safety measures in that setting. And the same is true with respect to various other contexts, such as pollution and public health. In these and other settings, we see extensive interventions in the name of promoting safety or well-being (whether all such interventions are genuinely motivated by such noble goals, and whether they are warranted in light of offsetting costs, are entirely different questions, of course).

So too with cybersecurity policy. Just like the building-safety context, there are an array of additional incentives that have been (or could be) brought to bear to incentivize greater investment in safeguards in the cybersecurity context. Indeed, many of the levers are the same. In both settings we find regulators who might dictate rules and standards, litigants who might sue, insurers who might set and enforce coverage conditions, and contract counterparties who might insist on certain actions as a condition of the deal.

These external incentive structures are not the only levers promoting improved defense on a systematic basis, however. Another involves efforts to promote sharing of information that might assist the cause of improved defense. And another tool is "pruning"—that is, legislation that repeals or carves out exceptions to existing laws that may be useful for other purposes, but that has the collateral impact of discouraging desirable defensive measures.

In the pages that follow in this subsection, we will survey each of these levers before concluding with a look at how the U.S. government organizes to promote better defense of (1) its *own* systems and (2) the systems associated with "critical infrastructure" owned by the private sector.

10. The Role of Regulators (I)

In theory, one way to promote better defense on a systematic basis is to create legally-binding rules or standards that oblige potential victims to take certain steps towards that end. Who might do that? In the first instance, of course, we should look to legislatures (both state and federal). In our separation-of-powers system, after all, that is the part of government meant to make laws. That said, anyone with even glancing familiarity with federal and state government over the past century will understand that a vast amount of rulemaking—for better or worse—stems not directly from legislation but rather from regulations promulgated by "administrative agencies."

A. About Federal Administrative Agencies

The federal government contains a large number of administrative agencies. Each has some particular field of subject-matter responsibility (the scope of which is defined by statute in most cases). Each typically performs many functions, but we are especially concerned with two core capacities.

First, rulemaking. An agency might have authority to promulgate legally-binding regulations (that is, to engage in "rulemaking") in furtherance of some goal specified by Congress. For example, Congress has given the Environmental Protection Agency authority to promulgate regulations to further the goals of the Clean Air Act in certain ways. There are a host of complex procedural

rules associated with agency rulemaking, but for now it is enough to know that this has been a common mode of creating law since the 20th century.

Second, enforcement. Congress sometimes authorizes an agency to initiate and pursue "enforcement" proceedings. The idea is that the agency may be tasked with investigating possible rule violations (whether a rule stated directly in a statute enacted by Congress, or a rule promulgated by an agency pursuant to authority delegated by Congress) and then initiating civil proceedings to enforce alleged violations. In some cases, the enforcement action might take the form of an ordinary civil suit, with the agency suing the alleged violator in federal court. But sometimes Congress empowers the agency also or instead to adjudicate the enforcement process internally (at least as an initial matter), with a litigation process involving an administrative law judge within the agency itself. Either way, the general idea is to secure a determination that someone violated a rule, producing a costly fine/damages, an order obliging (enjoining) the violator to take or cease some particular action(s), or both.

Like other forms of litigation, agency enforcement proceedings routinely result in settlements in which the alleged violator agrees to take or cease certain actions, with the possibility of more severe consequences later on if the party breaches that obligation.

Note that other parties in an industry may take note of the initiation and resolution of enforcement actions. Enforcement actions thus can cast a shadow—sometimes a very long shadow—impacting how other players decide to act.

> Bearing that in mind, can you make an argument that "enforcement" authority is itself a second form of rule-making authority?

B. There Is Not (Yet) An "EPA" For Cybersecurity

I mentioned the EPA above. It was created during the Nixon Administration at a time of mounting concern about the harmful effects of pollution. Over time, Congress has granted various rulemaking and enforcement authorities to the EPA in furtherance of this general mission. One might argue that mounting concern about inadequate cybersecurity warrants creation of a similar dedicated agency. It is important to grasp, however, that Congress so far has not taken that step. As we will see later in this unit, there *is* an important agency within the Department of Homeland Security known now (though not originally) as the Cybersecurity and Infrastructure Security Agency (CISA), but its role does not (yet) include rulemaking and enforcement authority encompassing the private sector.

All that said: just as nature abhors a physical vacuum, so too government bureaucracies seem to abhor perceived regulatory vacuums. There are existing administrative agencies with broad jurisdiction that have, in various ways, taken up at least some aspects of the cybersecurity regulatory challenge—sometimes because Congress has directed them expressly to do so, and sometimes because they have noted that their existing jurisdiction arguably encompasses certain aspects of cybersecurity and they have acted accordingly.

Your goal in light of all that: understand how certain pre-existing agencies have managed to participate in cybersecurity promotion. We'll start with the one you hear about the most in this space: the FTC.

C. The Federal Trade Commission ("FTC") and The FTC Act

For a very brief introduction to the FTC, read this.

> Based only on this overview, would you expect the FTC to have a role in setting or enforcing standards for cybersecurity? Why or why not?

One of the statutes the FTC is empowered to enforce is the Federal Trade Commission Act ("FTC Act"). The FTC has *not* engaged in any rulemaking relating to cybersecurity under this statute. Instead, it has focused on enforcement actions, based on the claim that some situations involving poor cybersecurity violated a rule set forth in the FTC Act itself. In fact, it has initiated more than 60 enforcement actions along these lines, and has touted the resulting body of cases as functioning, collectively, as a form of guidance to the private sector. In a moment I'll ask you to consider whether this level of enforcement, without tailored rulemaking, is desirable. But first you need to know just what they've been enforcing and how they've been doing it.

Goal #1: Understand what exactly the FTC Act prohibits. The answer is found in 15 U.S.C. 45(a)(1).

> **(1)** Unfair methods of competition in or affecting commerce, and unfair or deceptive acts or practices in or affecting commerce, are hereby declared unlawful.

There's an opening clause about unfair competition, and a second clause that refers alternatively both to "unfair" practices and "deceptive" practices that impact interstate commerce.

> Consider the following questions:
> - Can you explain how unfairness is different from deception?
> - Does either concept seem relevant to a situation in which some entity has poor cybersecurity?
> - In terms of clarity (and thus understanding on the part of those who must comply), how does this compare to the various subparts of the CFAA?

Goal #2: Understand how 15 U.S.C. 45(n) limits one (but not both) of those two prohibitions.

> **(n)** **Standard of proof; public policy considerations**
>
> The Commission shall have no authority under this section or section 57a of this title to declare unlawful an act or practice on the grounds that such act or practice is unfair unless the act or practice causes or is likely to cause substantial injury to consumers which is not reasonably avoidable by consumers themselves and not outweighed by countervailing benefits to consumers or to competition. In determining whether an act or practice is unfair, the Commission may consider established public policies as evidence to be considered with all other evidence. Such public policy considerations may not serve as a primary basis for such determination.

> Which one is impacted, and is the impact likely relevant for a cybersecurity situation?

Goal #3: Understand whether there are significant limits with respect to who has to care about the FTC Act. Read 15 U.S.C. 45(a)(2).

> **(2)** The Commission is hereby empowered and directed to prevent persons, partnerships, or corporations, except banks, savings and loan institutions described in section 57a(f)(3) of this title, Federal credit unions described in section 57a(f)(4) of this title, common carriers subject to the Acts to regulate commerce, air carriers and foreign air carriers subject to part A of subtitle VII of title 49, and persons, partnerships, or corporations insofar as they are subject to the Packers and Stockyards Act, 1921, as amended [7 U.S.C. 181 et seq.], except as provided in section 406(b) of said Act [7 U.S.C. 227(b)], from using unfair methods of competition in or affecting commerce and unfair or deceptive acts or practices in or affecting commerce.

Does subsection 45(a)(2) encompass everyone?

Goal #4: Understand how the FTC goes about enforcing the FTC Act. It has two available options. Read 45(b) and 45(m).

> **(b) Proceeding by Commission; modifying and setting aside orders**
>
> Whenever the Commission shall have reason to believe that any such person, partnership, or corporation has been or is using any unfair method of competition or unfair or deceptive act or practice in or affecting commerce, and if it shall appear to the Commission that a proceeding by it in respect thereof would be to the interest of the public, it shall issue and serve upon such person, partnership, or corporation a complaint stating its charges in that respect and containing a notice of a hearing upon a day and at a place therein fixed at least thirty days after the service of said complaint. The person, partnership, or corporation so complained of shall have the right to appear at the place and time so fixed and show cause why an order should not be entered by the Commission requiring such person, partnership, or corporation to cease and desist from the violation of the law so charged in said complaint. Any person, partnership, or corporation may make application, and upon good cause shown may be allowed by the Commission to intervene and appear in said proceeding by counsel or in person. The testimony in any such proceeding shall be reduced to writing and filed in the office of the Commission. If upon such hearing the Commission shall be of the opinion that the method of competition or the act or practice in question is prohibited by this subchapter, it shall make a report in writing in which it shall state its findings as to the facts and shall issue and cause to be served on such person, partnership, or corporation an order requiring such person, partnership, or corporation to cease and desist from using such method of competition or such act or practice. Until the expiration of the time allowed for filing a petition for review, if no such petition has been duly filed within such time, or, if a petition for review has been filed within such time then until the record in the proceeding has been filed in a court of appeals of the United States, as hereinafter provided, the Commission may at any time, upon such notice and in such manner as it shall deem proper, modify or set aside, in whole or in part, any report or any order made or issued by it under this section. After the expiration of the time allowed for filing a petition for review, if no such

petition has been duly filed within such time, the Commission may at any time, after notice and opportunity for hearing, reopen and alter, modify, or set aside, in whole or in part any report or order made or issued by it under this section, whenever in the opinion of the Commission conditions of fact or of law have so changed as to require such action or if the public interest shall so require, except that (1) the said person, partnership, or corporation may, within sixty days after service upon him or it of said report or order entered after such a reopening, obtain a review thereof in the appropriate court of appeals of the United States, in the manner provided in subsection (c) of this section; and (2) in the case of an order, the Commission shall reopen any such order to consider whether such order (including any affirmative relief provision contained in such order) should be altered, modified, or set aside, in whole or in part, if the person, partnership, or corporation involved files a request with the Commission which makes a satisfactory showing that changed conditions of law or fact require such order to be altered, modified, or set aside, in whole or in part. The Commission shall determine whether to alter, modify, or set aside any order of the Commission in response to a request made by a person, partnership, or corporation under paragraph [1] (2) not later than 120 days after the date of the filing of such request.

(m) Civil actions for recovery of penalties for knowing violations of rules and cease and desist orders respecting unfair or deceptive acts or practices; jurisdiction; maximum amount of penalties; continuing violations; de novo determinations; compromise or settlement procedure

(1)

(A) The Commission may commence a civil action to recover a civil penalty in a district court of the United States against any person, partnership, or corporation which violates any rule under this subchapter respecting unfair or deceptive acts or practices (other than an interpretive rule or a rule violation of which the Commission has provided is not an unfair or deceptive act or practice in violation of subsection (a)(1)) with actual knowledge or knowledge fairly implied on the basis of objective circumstances that such act is unfair or deceptive and is prohibited by such rule. In such action, such person, partnership, or corporation shall be liable for a civil penalty of not more than $10,000 for each violation.

(B) If the Commission determines in a proceeding under subsection (b) that any act or practice is unfair or deceptive, and issues a final cease and desist order, other than a consent order, with respect to such act or practice, then the Commission may commence a civil action to obtain a civil penalty in a district court of the United States against any person, partnership, or corporation which engages in such act or practice—

(1) after such cease and desist order becomes final (whether or not such person, partnership, or corporation was subject to such cease and desist order), and

(2) with actual knowledge that such act or practice is unfair or deceptive and is unlawful under subsection (a)(1) of this section.

> In such action, such person, partnership, or corporation shall be liable for a civil penalty of not more than $10,000 for each violation.
>
> **(C)** In the case of a violation through continuing failure to comply with a rule or with subsection (a)(1), each day of continuance of such failure shall be treated as a separate violation, for purposes of subparagraphs (A) and (B). In determining the amount of such a civil penalty, the court shall take into account the degree of culpability, any history of prior such conduct, ability to pay, effect on ability to continue to do business, and such other matters as justice may require.
>
> **(2)** If the cease and desist order establishing that the act or practice is unfair or deceptive was not issued against the defendant in a civil penalty action under paragraph (1)(B) the issues of fact in such action against such defendant shall be tried de novo. Upon request of any party to such an action against such defendant, the court shall also review the determination of law made by the Commission in the proceeding under subsection (b) that the act or practice which was the subject of such proceeding constituted an unfair or deceptive act or practice in violation of subsection (a).
> **(3)** The Commission may compromise or settle any action for a civil penalty if such compromise or settlement is accompanied by a public statement of its reasons and is approved by the court.

Can you explain the difference between the two procedures in terms of who decides whether the FTC's allegation is correct? In terms of what remedies appear to be available if the FTC wins?

Case Study #1: Uber

Read this complaint filed by the FTC.

Consider the following questions:

- Which enforcement path did the FTC use in this case?
- In what way(s) did Uber allegedly violate Section 45(a)?
- Assuming all the allegations to be true, would you agree with the FTC that this violates the statute?

Eventually Uber settled with the FTC, but then in April 2018, the FTC announced it had reopened the case due to Uber's failure to disclose that, during the pendency of the case at the earlier stage, Uber had experienced another data breach. This led to a revised settlement agreement.

Scan the document to get a sense of the commitments FTC extracted from Uber.
- What are those commitments, and was this a good outcome?

The FTC was not the only problem Uber faced in connection with these events. A number of state Attorneys General decided to team up and sue Uber together, based on various state data-breach liability laws we will examine in the next class. For now, it is enough to note that Uber faced this massive and well-resourced lawsuit at the same time that it faced the FTC's renewed enforcement action.

Consider how the pendency of parallel litigation might matter. Oh, by the way, Uber and the states have just announced (9/26/18) that they are settling for $148 million.

Case Study #2: Wyndham Hotels

Here's an example of the FTC suing in federal court. Read this note summarizing the litigation involving Wyndham Hotels.

Consider the following questions:
- What was Wyndham's argument about the propriety of suing them under the "unfairness" prong of Section 45(a)(1)?
  - How did the court rule on that point, and do you agree?
- What was Wyndham's second argument, concerning "fair notice"?
  - How did the court rule on that one, and do you agree?

Note: Wyndham and the FTC settled later, with Wyndham agreeing to take on a variety of security-focused practices (as well as annual audits) for the next 20 years.

Case Study #3: LabMD

This was a remarkable case in many respects. I won't summarize it here, but rather will ask you to read this short overview from Prof. Dan Solove.

Can you summarize how the outcome in *LabMD* compares to *Wyndham*?

11. The Role of Regulators (II)

A. The FTC and the Gramm-Leach-Bliley Act

As it happens, the FTC also has authority to enforce other statutes, and in some cases to promulgate regulations relating to them. One such statute is the Gramm-Leach-Bliley Act (the "GLB Act"), which, among other things, concerns the protection of customer data by financial institutions.

The FTC has promulgated a set of regulations on that issue, known collectively as the "Safeguards Rule" (found in 16 Code of Federal Regulations Part 314). Skim the text of Part 314 and then read this FTC-written overview.

> Who does this govern, and (at a general level) what does it require them to do?

For a recent illustration of the Safeguards Rule in action, read pp.3-5 of this action the FTC pursued against TaxSlayer. I won't have questions for you about this one; it's just an illustration.

B. A Quick Look at Other Federal Regulators

There are other federal regulators involved in cybersecurity, besides the FTC. We will not go into anything like the same level of detail with them, but you should have at least a glancing familiarity with the roles some of them play.

> For each example below, identify the substantive standard that the agency appears to be enforcing.

Read here for an example involving the Securities & Exchange Commission ("SEC").

Skim this (just glance through the first dozen pages) for an example involving the Federal Communications Commission ("FCC").

Note that medical devices obviously raise especially acute cybersecurity concerns, particularly when the device in question can be accessed remotely and is capable of causing significant harm. I'm not assigning you anything relating to the Food and Drug Administration (the "FDA"), but if you are interested in going deeper on this topic: do some searching to see if you can determine whether the FDA has gotten involved with cybersecurity regulations or enforcement.

C. Don't Forget the State Regulators and Foreign Regulators

Why should all the fun be left to federal regulators? Of course, it isn't. I want you to be generally aware of various ways in which other regulators become involved, though I'm not going to ask questions about this in class or hold you accountable on the exam for it. This is just for your general awareness.

This example recently took effect in New York in relation to the financial services industry. And here's a short piece explaining how various European regulators responded to the same Uber breach we noted above.

12. Private Lawsuits (I)

A. The Big Picture: Who Are the Injured Parties Who Might Become Plaintiffs?

Consider the following depiction of the chain of actors involved in a common type of cybersecurity incident:

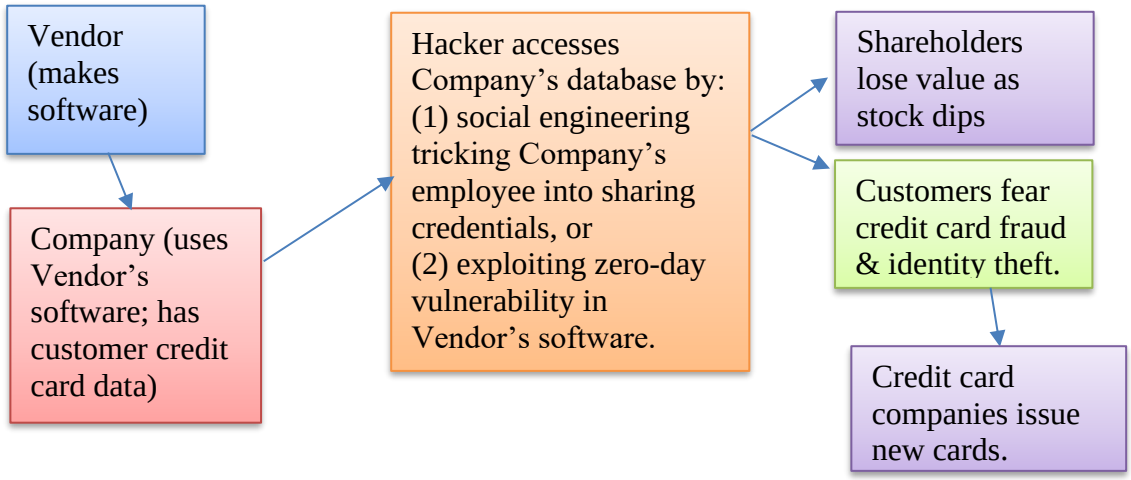

Looking at the chart above, consider the following questions:

- Who would you characterize as a victim?
- One might expect that if anyone is to be sued for damages, it would be the hacker. That rarely occurs, however. Why might that be?

Usually the vendor is not sued either. That's to be expected if the hacker breached the company's security via social engineering; that's a failure on the part of the employee(s) and perhaps the company, but not the vendor. But what if the breach was the result of a vulnerable in the vendor's software? Read this article.

Make a list of the obstacles to suing software providers.

- Consider the policies that might be served by each obstacle. Would you change any of them?

This leaves the option of suing the company that actually suffered the breach. Our goal now is to understand the major form of liability that companies in this situation might face.

B. Tort Liability (Liability For Lack of Due Care)

In our legal system, we use the word "tort" to refer to situations in which the law authorizes suits for damages based on harmful actions/omissions. Some torts are "common law" causes of action, meaning that the courts have recognized a right to sue in a particular situation even without a statute calling for recognition of that tort. This is the traditional and most familiar kind of tort. Examples include negligence and battery. But legislatures can create torts by statute, too, if they wish.

There are many kinds of torts. One batch, called "intentional torts," involves purposefully harmful conduct. That's not our concern here, for we are assuming that companies do not intend to be breached. So that leaves us with *unintentional* harms. In that situation, the tort system can take either of two approaches. First, it can make someone strictly liable for all harms they cause. Second, it can make them liable only for harms that result from lack of adequate care—what we commonly call "negligence." The strict-liability approach is relatively rare, and usually confined to ultra-hazardous activities. For companies that may have inadequate information security, the important question is negligence.

As all law students learn during their first year, negligence makes a defendant liable in damages where four conditions are met: (1) the defendant owed a duty of care, (2) the defendant breached that duty, (3) the plaintiff suffered a legally recognizable harm, and (4) the breach was the proximate (reasonably foreseeable) cause of that harm.

> Pause now to consider how, in the scenario above, the company's customers *might* have a negligence claim against (a) the company or (b) the vendor.

Plaintiffs have an uphill battle in bringing negligence claims against companies that suffer a data breach, because each of the elements of a traditional negligence action are difficult to show in the cybersecurity context. For example, even if a court assumes that a business has a duty to protect its customers' data, plaintiffs must show the business's cybersecurity practices were so substandard that it actually breached that duty of care—and this requires comparing those practices to some standard of what "reasonable" cybersecurity practices look like.

Next, plaintiffs must show that they suffered some harm as a result of the breach. But how? For plaintiffs whose data has been exposed, it can be difficult to identify the amount of harm they actually suffered unless that data has been used against them. But even for plaintiffs who can show they suffered identity theft or some other harm, how can they prove it was the result of this particular breach, and not some other breach of their data?

The Supreme Court recently denied certiorari in two cases that would have addressed the problems plaintiffs face in obtaining class certification in class-action suits. Here is what Joseph Marks at the Washington Post had to say about it in December 2018:

> Two class-action lawsuits that could come before the Supreme Court this term seek to determine just how bad a cybersecurity lapse must be before customers can sue.
>
> In both cases, federal appeals court judges formally approved lawsuits by thousands of consumers who want to collectively sue major companies for cybersecurity failures — even though the customers couldn't prove they'd suffered any direct financial harm from the companies' digital negligence.

The companies are asking the high court to overturn the lower court decisions allowing the lawsuits. They argue that customers must suffer some concrete financial or physical harm before they can demand compensation for a data breach or for hackable vulnerabilities discovered in their products.

Flexible tech solutions can prepare your business for nearly any challenge, but first you need a game plan.

Consumers, however, contend that setting such strict standards would give negligent companies a pass for not sufficiently protecting their products and data.

If the Supreme Court rules on either case, it could fundamentally reshape the responsibility the private sector has over the security of Internet-connected products that could endanger consumers' privacy or even their lives in the case of things like cars and medical devices.

If the court sets a high bar for consumers to sue, it could prompt companies to play fast and loose with their data. If that standard is too low, however, it may deter companies from sharing information about newfound computer bugs or investing in new technologies out of fear they'll be on the hook for legal damages.

"You've potentially jacked way up the monetary costs from a vulnerability that's disclosed down the road," Megan L. Brown, an attorney with Wiley Rein who deals in complex litigation and technology, told me. "That may affect a company's risk calculation and make them not do some things."

The first class action suit was sparked after a viral 2015 Wired article describing how two security researchers hacked through the entertainment system in a Jeep Cherokee to kill the brakes — all while the Wired reporter was driving the vehicle at 70 mph through downtown St. Louis.

After the article, Chrysler mailed 1.4 million vehicle owners a USB stick with software to fix the vulnerability, and there's no evidence malicious hackers ever exploited it. Jeep owners point to the hack, however, as evidence that their vehicles are "excessively vulnerable" and say they should get some money back, according to Chrysler's petition to the high court.

The issue is particularly complicated because cybersecurity experts warn there's no way to ensure any system is 100 percent digitally secure.

Even major digital consumer products such as Microsoft's Office suite or Apple's iPhone aren't invulnerable. Security researchers find hackable vulnerabilities in those products every week. The most mature and cyber-sensitive companies, however, usually manage to find and patch the most dangerous vulnerabilities before malicious hackers exploit them.

If the Jeep plaintiffs are successful, "it opens the door to litigation of all stripes and flavors over any consumer product that connects to the Internet," Chrysler attorney Thomas H. Dupree Jr. told me. "Any product, hypothetically, can be hacked and any plaintiff can hire a lawyer who says, 'In my opinion the product has inadequate cybersecurity even though it hasn't been breached.' "

In the second case, the online retailer Zappos did suffer a malicious breach of a database containing customers' information, including names, contact information and possibly credit card data. The company says, however, that there's no evidence the hackers used that data to impersonate customers or to make phony credit card charges.

The customers dispute that characterization, however, and say hackers used their information to hack other accounts.

The Zappos and Jeep cases are being litigated at the U.S. District Court level while the lower courts wait to learn whether the Supreme Court will hear the cases. Neither case has moved substantially past the questions of whether the plaintiffs have standing to sue and who should be included in the plaintiffs' class.

The high court held a conference on the Zappos case this month and is scheduled to meet on the Jeep case Jan. 4. The court probably will decide whether to grant hearings in the cases in January.

Meanwhile, industry groups are worried about the potential implications. Trade associations like the U.S. Chamber of Commerce, the National Association of Manufacturers and CTI, the wireless association, have filed friend-of-the-court briefs supporting the companies.

They have reason to be concerned if the high court does take the cases. Lawsuits where there's much clearer harm from a data breach have resulted in multimillion-dollar settlements. Target, for example, paid $18.5 million to settle cases brought by state attorneys general over its 2013 breach of credit and debit card information for at least 40 million customers and personal information about many more. The retailer is trying to conclude a $10 million settlement on a consumer class action stemming from that breach.

In other cases, assessing whether hack victims have suffered harm can be far more difficult.

The U.S. Court of Appeals for the D.C. Circuit, for example, is mulling a case filed by federal employee unions over the 2015 Office of Personnel Management data breach. Most cyber experts believe that breach was launched by Chinese government hackers who want to use the data for blackmail or other espionage.

That means it's unlikely the data stolen from more than 21 million current and former federal employees and their families will be used for identity theft or to make phony credit card charges. The stolen data, however, included extremely personal background check information, including lengthy questionnaires about finances, housing, family relationships and drug and alcohol use.

An appeals court judge said during a Nov. 2 hearing that the government faced an "uphill battle" arguing the plaintiffs didn't have grounds to sue in that case, as reported by GovExec.

The harm caused by that kind of hack is far different from the nebulous damage caused by a breach involving only information such as names and addresses, Joe Hall, chief technologist at the Center for Democracy and Technology think tank, told me.

> "Trivial harm should not be something that keeps us from building wonderful things," Hall said. "But we really need to find a way to articulate harms that are not economic but really affect people's ability to trust each other or to participate in the world."

So, what happened next? The Supreme Court has declined to hear either case, leaving the Court of Appeals rulings in place.

Next, consider the following case study, arising out of the infamous data breach at the credit agency Equifax. Credit-reporting agencies like Equifax are a particularly appealing target, in light of the vast volume and sensitive nature of the information they collect. As you might imagine, news of the breach made headlines, and many lawsuits followed. Eventually, a group of plaintiffs joined to seek "class action" status—that is, to represent and seek recovery not just on their own behalf but also on behalf of every other similarly-situated person. Here is an account of that suit, from Anne Bucher:

> Equifax Inc. has been hit with a class action lawsuit filed on behalf of individuals in all 50 states and the District of Columbia over the massive Equifax data breach that was announced in September.
>
> A massive cybersecurity incident reportedly exposed the sensitive personal identification and financial information of more than 145 million people—nearly half of the U.S population.
>
> "This case concerns the largest data breach involving personal and financial information in American history," the Equifax class action lawsuit says. "Equifax, one of the three credit reporting companies used by thousands of businesses to assess the credit worthiness of customers and prospective customers, failed spectacularly in protecting that data."
>
> As a result, sensitive consumer information such as names, birth dates, Social Security numbers, drivers' license numbers, address history and other financial information was allegedly accessed by unauthorized parties from May through June.
>
> The Equifax class action lawsuit claims the damage caused by the Equifax data breach was compounded by the credit reporting agency's "egregious cybersecurity failures before, during, and after the breach," including: failing to employ an available security patch, not recognizing the data breach for more than three months, failing to implement security measures after the data breach, and failing to timely inform the public about the massive data breach.
>
> The plaintiffs also challenge Equifax's response to the data breach, including the confusing emails it sent to consumers about whether their information was compromised, and creating a credit monitoring service with a confusing message about whether users would be bound by an arbitration clause, and sending consumers a wrong link to have their credit frozen.
>
> … This new 323-page Equifax class action lawsuit seeks to consolidate dozens of data breach lawsuits that are currently pending throughout the country. At least 150 data breach lawsuits have been filed nationwide against Equifax so far.
>
> The plaintiffs include individuals from all 50 states and the District of Columbia who allege they have experienced fraud after their sensitive personal information was compromised in the data breach.

> The Equifax class action lawsuit asserts claims for violation of the Fair Credit Reporting Act, negligence, negligence *per se*, bailment, unjust enrichment, and violation of state consumer protection and/or privacy laws. The plaintiffs are seeking monetary damages, declaratory and injunctive relief, and other remedies allowed under the applicable laws.

The plaintiffs in this case sought "class action" status—that is, they sought approval from the court to assert not just their own claims but those of all other similarly-situated persons.

Consider the pros and cons of class-action status from the point of view of the defendant, the plaintiff, and the plaintiff's attorneys (bearing in mind that the plaintiff's attorneys most likely are being compensated on a "contingent fee" basis—meaning that they will receive a percentage of the eventual recovery, if any).

It can be hard to prove what the duty of care is, let alone that a breach of it occurred. Read this CNN account of the Equifax hack:

> How did it happen? Much is still unknown. But it came down to a flaw in a tool designed to build web applications, the company said in a press release this week. And Equifax admitted it was aware of the security flaw a full two months before the company says hackers first gained accessed to its data.
>
> Some of the information hackers had access to includes names, Social Security numbers, birth dates, addresses and some driver's license numbers.
>
> The tool is called Apache Struts, and it's used by many large businesses and government organizations. Equifax used it to support its online dispute portal -- where Equifax (EFX) customers go to log issues with their credit reports. The flaw allowed hackers to take control of a website.
>
> A cybersecurity arm of the U.S. Department of Homeland Security, US-CERT, "identified and disclosed" the Apache Struts flaw in March, Equifax said in a statement.
>
> And the company's security department "was aware of this vulnerability at that time, and took efforts to identify and to patch any vulnerable systems."
>
> Yet, according to the company, hackers exploited the flaw months later.
>
> Equifax has said it discovered the data breach on July 29. On Friday, it said it waited until it "observed additional suspicious activity" a day later to take the affected web application offline.
>
> And on August 2 Equifax contacted Mandiant, a professional cybersecurity firm, to help the company assess what data had been compromised.
>
> With help from Mandiant, Equifax was able to determine a series of breaches had occurred from May 13 through July 30, the company said.
>
> Patching software at big corporations with many machines does take time. They must first identify the vulnerability, then implement and test the patch to make sure it doesn't break anything before making it public.

> However, security experts say Equifax should have moved faster.
>
> "There's really no excuse whether it's a difficult patch or not, for an organization of that size with that kind of magnitude of data," said Jon Hendren, director of strategy at security firm UpGuard. "When you're a big organization like that, it's a systemic failure of process and the blame goes straight to the top."
>
> Equifax has also been widely criticized for waiting more than a month to alert its customers and shareholders about the hack.
>
> On Friday, the company announced its chief information officer and chief security officer are "retiring."

Consider the following questions:

- Assume that is all accurate. Do you feel Equifax breached a duty of care?
- What, exactly, is the duty?
- If you think Equifax breached its duty, what specifically would you say it should have done? Can you explain how your understanding of Equifax's duty is different from saying that Equifax is simply strictly liable for any breach that might occur?
- What do you suppose Equifax would argue in response?

Proving the existence of a duty of care, and a breach, can be difficult. But even if proven, that's not enough. The plaintiff also must prove that the breach was the proximate cause of damage. This can be a huge challenge in the context of data breaches.

Read this from the Electronic Frontier Foundation for a critical perspective on that problem.

- Do you agree with EFF?

So, what happened in the Equifax case? See here.

Consider the following questions:

- Why did the company settle?
- Is this a good outcome from the point of view of the larger public interest(s) at stake?

You might be wondering whether the FTC and other such regulators also took an interest in the Equifax case. Oh yes, they certainly did. See here.

Consider these questions:

- What do you suppose the FTC's theory for an enforcement action actually was?
- Does the combination of both private sector litigation and regulatory enforcement actions seem proper to you?

C. Statutory Interventions to Clarify Liability

Note that legislatures can intervene to clarify or modify the scope of tort-style liability if they wish, making it harder or easier to sue. California, for example, recently did exactly that. Section 11 of the California Consumer Privacy Act of 2018 (codified at Cal. Civil Code Section 1798.150) provides that certain companies doing business in California are subject to civil suit for injunctive relief or damages in the amount of the plaintiff's actual damages or else "statutory damages" (that is, a pre-set penalty determined by the statute) in the range of $100-$750 per consumer per incident, whichever turns out to be higher, if "nonencrypted or nonredacted personal information . . . is subject to an unauthorized access and exfiltration, theft, or disclosure as a result of the business's violation **of the duty to implement and maintain reasonable security procedures and practices appropriate to the nature of the information.…**"

Consider these questions:

- How does this statutory standard compare to common-law negligence?
- In light of how the statute describes the category of information subject to the statute's protection, what advice would you give to a company that has personal information to protect?

13. Private Lawsuits (II); Insurance & Contract Terms

A. Contract Liability (For Failing to Live Up to Security-Related Promises)

In some settings, the company that suffered a breach will have made security-related representations in a contract of some kind. A professional-services firm, for example, might include such representations in the "engagement letter" that serves as the contract between the firm and its clients (sophisticated clients increasingly will insist upon this). In other settings, there may be terms-of-service that govern a customer's or user's relationship to a company (particularly, but not only, where customers or users interact with the company via an app). These too may contain security-related representations. These and other examples create the possibility of a breach-of-contract lawsuit in the event of a data breach where the breach arguably shows that the company failed to live up to its promises.

Of course, companies make some representations in settings that do not count as part of a contract with a customer or user. For example, a company may make statements in advertisements or on their websites, including statements about care they take to protect customer and user data.

Consider how this illustrates the difference between a breach of contract claim and an FTC enforcement action based on deceptive advertising.

Case Study: Anthem

In February 2015, the health insurance company Anthem announced that its security had been breached and that a massive amount of personally identifiable information about patients had been exposed. This led to massive litigation, based on a variety of claims, including breach-of-contract claims. Anthem tried to have these claims dismissed but was only partially successful. On one hand, it *did* succeed in having the breach-of-contract claims dismissed, on the ground that the promises it made regarding customer privacy on its website's "privacy statement" and in certain mailings to customers were not actually part of a contract with customers. On the other hand, the court refused to dismiss a separate cause of action, under California-state law, for deceptive advertising.

Having failed to get the whole case dismissed, Anthem eventually settled. Read about it here and consider the following questions:

- What did the plaintiffs receive? What did the plaintiff's *attorneys* receive?
- How do you feel about this result?

Also read this more recent notice about the settlement.

By this point, you surely are asking yourself: Didn't the regulators also get in on the Anthem action? Of course they did (just like they also got in on the Equifax action).

Because Anthem was in the insurance business, a California state insurance regulatory agency conducted an investigation. It took a considerably more-sympathetic view of the matter, emphasizing state-actor responsibility for the breach and suggesting that there was little chance Anthem could have long withstood the persistent efforts of a nation state to capture this data.

Is this inconsistent with the civil litigation result? If so, is that a policy problem, and what might be done about it?

As for the FTC, It did not get involved. But because Anthem was involved in health matters, another federal agency did: the Department of Health and Human Services ("HHS"). For a quick glance at HHS's role in this space, read this.

B. Liability for Inadequate Disclosure of a Data Breach

The possibility of being sued for inadequate care, failing to live up to a promise, or deceptive advertising all loom large when a company learns that it has been breached and that customer or user data has been exposed. And that in turn creates an incentive for the company leadership to proceed very carefully—and thus very slowly—in letting anyone know that the breach has occurred. But there is a powerful consideration pushing in the exact opposite direction: all states have laws compelling companies to disclose such breaches, and to do so rapidly.

Can you summarize the competing public-policy interests implicated by such laws, and how you assess the balance between them?

Needless to say, we are not going to look at all the different state-breach-disclosure laws. Instead, we'll look at Texas law as an example.

Read Texas Business and Commerce Code Section 521.053(b), and consider these questions:

- How certain must the entity be that a breach occurred?
- Precisely how quickly must the disclosure be made as a default matter?
- How does Section 521.053(d) potentially change that timeline?

The Texas Attorney General is responsible for filing civil suits to enforce Section 521.053, but note that a Section 521.053 violation may also constitute a deceptive act for purposes of a private civil suit under the Texas Deceptive Trade Practices Act.

As you might imagine, the patchwork quilt of disclosure laws around the various states (as well as some cities) has led some to argue for Congress to impose a uniform national approach.

What are the pros and cons of one national standard?

Don't forget: Some companies will be subject to foreign jurisdictions as well, and these may be more demanding than American disclosure requirements. The European Union's much discussed "General Data Protection Regulation" ("GDPR") is a particularly significant example you should be aware of (though we are not studying its particular requirements in this class).

C. Shareholder Derivative Actions

Recall that in our generic scenario above, we noted the possibility that Company B has shareholders who might sue it, assuming share prices dropped once the breach became public. Such "shareholder derivative actions" arise frequently with publicly traded companies when those companies experience any sort of significant reversal that might be attributed to bad decision-making by the company's officers or board of directors. Read this article for a fine overview of how such suits have fared in some of the most well-known data-breach cases.

D. Insurance

In any context in which entities and individuals can anticipate suffering a loss—whether the loss be from damage to possessions or a person, or from at least some forms of legal liability—there is a strong incentive to protect against the anticipated loss by purchasing an insurance policy. Because insurers usually (though definitely not always) are at liberty to determine which sorts of risks they will insure against, and subject to which conditions, the insurance industry in general is in a powerful position to nudge or even compel certain behaviors (just think of the incentives for safe driving that car insurance does or might generate). And thus, insurance has an important role to play in relation to the general challenge of encouraging potential victims to engage in better defense.

For a handy and accessible (and brief) introduction to the emerging cybersecurity insurance market, read this CEIP report, "*Addressing the Private Sector Cybersecurity Predicament: The Indispensable Role of Insurance*" (Nov. 2018), by Ariel Levite, Scott Kannry, and Wyatt Hoffman.

E. Contract Terms

When an insurer insists upon certain cybersecurity-related measures as a condition of coverage, that is an example of using contact leverage. And insurers are by no means the only ones who might choose to use that leverage for this purpose. In fact, private sector actors increasingly often insist upon certain steps as a contract condition, particularly in contexts involving professional services (*e.g.*, with law firms) or supply chains. While flowing from disaggregated individual self-interest rather than top-down government policy, this phenomenon has considerable potential to generate systematic improvements in the overall level of cybersecurity preparedness.

That said, there are times when the government itself can use contract leverage toward this same end—that is, not just to protect itself (we'll see that later, when we reach the subunit on government efforts to improve its own defenses) but also to have a larger pro-defense effect.

Here's an example involving an attempt by the government to leverage contracting power to keep certain private entities from using the antivirus products and other services of Kaspersky Lab (a Russia-based AV vendor that once had a substantial share of the US market, and still has a large global presence).

> Can you identify limits to the utility of this approach, based on this example? What are the pros and cons?

14. Facilitating Better Defense through Information-Sharing

Consider for a moment the idea of "public health" policy—that is, the important set of goals associated with minimizing and mitigating disease and injury on a systematic basis. As with cybersecurity, a wide variety of both public and private actors seek to promote public health. And as with cybersecurity, those actors employ many different tools. Some of these tools are familiar to us from the cybersecurity setting, in fact; regulations, for example, play a key role in promoting public health. And so, we might take the comparison a bit further, using the familiar public-health model as a way to appreciate another parallel with cybersecurity policy: facilitation of information-sharing as a means to promote safety.

Some amount of information-sharing will occur without external interventions by public or private-sector actors, of course. But the "natural" level of such sharing may be suboptimal (too limited, too slow, too confusing, too inaccurate, etc.) for any number of reasons. And if that is the case, then there is an argument for *someone* to take steps to improve things from that baseline level. In the public-health setting, we find governmental and international entities like the Centers for Disease Control and the World Health Organization which do this sort of thing, as well as an array of private sector actors (including both non-profit and for-profit entities). In various ways, these actors develop, curate, and disseminate a wide array of useful information, ranging from highly technical information to "best practices" and procedures.

So too with cybersecurity. That is, an important part of cybersecurity policy is the facilitation of information-sharing intended to facilitate and promote defense. Our goal in the next two classes is to become better acquainted with some of the key information-sharing actors and activities, and then to examine an important piece of legislation that attempts to prune away legal disincentives to information sharing.

A. What Information Might We Want to Share?

There are many possibilities. Let's start with the sharing of "threat intelligence." Note that this is a term of convenience (the meaning of which can vary from audience to audience) rather than a term of art (with a well-settled meaning and scope). That caveat aside, what is the general idea?

For starters, threat intelligence includes "Indicators of Compromise" (aka "IOCs"). This is a shorthand for the idea that there are technical signatures and other objective indicia that might be useful as a basis for detecting unauthorized activity in a system. This could be, for example, a signature element of an exploit, the fact that a system has a particularly important vuln, the URL of a known botnet command-and-control server, and so forth; the common idea is that the information is something one can search for in a scalable way, and that it is correlated with malicious activity.

For at least some people, threat intelligence has a broader scope than just IOCs. That is, the phrase might also encompass all sorts of additional, useful information, some of it technical (for example, details about a newly discovered vulnerability or newly circulated patch) and some of it otherwise (for example, information about the capabilities, motives, intentions, or characteristic tactics, techniques, and procedures of potentially hostile entities or individuals).

Be prepared to identify and distinguish these categories, and to explain why sharing information of each kind can be helpful for improving defense.

Threat intelligence is not the only relevant category of information we might hope to spread as a means to boost cybersecurity on a widespread basis. We might also want to spread knowledge of what we might call "effective practices" or "best practices."

The idea here is straightforward: we want as many individuals and organizations as possible to at least be *aware* of optimal policies, practices, techniques, or procedures, just as we do in the public-health setting (just think of the efforts made to encourage people to wash their hands, cover their sneezes, get flu shots, quit smoking, etc.). Sometimes we go further, of course, actually compelling adoption of such practices (just think of the FTC's current attempt to revise the G-L-B Act Safeguards Rule in order to require, among other things, the use of multi-factor authentication in some circumstances). But here we are talking about the less intrusive step of simply increasing the chance that someone will adopt such measures voluntarily.

This raises a question: If a practice is desirable, why not always go with mandatory rules rather than merely helpful advisories? First, sometimes the case for mandating use of a particular measure just isn't clear, as when the measure entails offsetting costs that might in some cases outweigh the benefits. In such circumstances, promoting awareness of the measure helps individuals and organizations make their own judgment about that cost-benefit tradeoff.
Second, even if the measure might be worthy of a mandate, there may be any number of reasons why it is not currently possible to get Congress or a relevant agency to take that particular step. Merely sharing information about the measure normally will not require anything like the same degree of political will. Indeed, as we will see below, it may not require government intervention at all.

Be able to explain how mandatory rules and informational advisories relate to each other, and why the former is generally more difficult to employ than the latter.

B. Why is Information-Sharing Difficult, in Theory?

Information-sharing can be conducted on a government-to-government, government-to-private, private-to-government, and private-to-private basis. Each dynamic presents its own challenges, giving rise to the possibility that we might not get sufficient sharing, on a systematic basis, without interventions.

Consider these questions:

- Why might one government entity be reluctant to share information with another?
- Why might the government be reluctant to share with the private sector?
- Why might the private sector be reluctant to share with the government?
- Why might one private sector entity resist sharing with another?
- Do your answers depend on which type of information is at issue?

C. Key Information-Sharing Institutions

Let's now have a look at some of the institutions and systems involved in promoting information-sharing relating to cybersecurity. What follows below is *not* meant to be comprehensive, by any stretch. That said, this selective survey will give you a feel for some of the key systems and institutional players, while underscoring that much of the work under this heading is carried out through the private sector.

Security Vendors:

We will start by observing that there is a large and dynamic "security vendor" community consisting of an array of for-profit businesses that provide various cybersecurity-related products and services. Of course, one must pay to get the benefits of these products and services in most circumstances, so most of the security benefits that vendors provide do not properly qualify for discussion under the heading of information-sharing. But vendors frequently do provide certain security-related information on a free and open basis, as when companies publish special "reports" identifying newly discovered malware campaigns.

Can you explain why a for-profit enterprise would publish such a report?

VirusTotal:

VirusTotal is a company founded in Spain in 2004, later acquired by Google, and now operated by Chronicle (a fellow subsidiary of Google's parent entity, Alphabet). The basic idea is simple: VirusTotal aggregates the capabilities of a large number of antivirus and URL/domain blacklisting companies, and provides a user-friendly way to check files against the sum of their capabilities. The value proposition for users is obvious enough, but what's in it for all those companies? Well, they get to see the results too, and thus can improve their own individual capabilities. It would not work if all those companies still competed (as they used to) primarily on the basis of having exclusive, propriety databases of threat indicators, but these days the bigger companies mostly embrace a model of open sharing of such information while competing based on other considerations. That's a big win for the public at large, and VirusTotal is a notable mechanism for transmitting the resulting benefits in a user-friendly way.

Can you describe how this shift is analogous to public health and pharmaceuticals?

For a more detailed explanation of what one can do with VirusTotal, and how the community-good model of information-sharing about threat indicators can be used as an instrument of statecraft, read the following excerpt from a 2018 article at Motherboard:

> This week, US Cyber Command (CYBERCOM), a part of the military tasked with hacking and cybersecurity focused missions, started publicly releasing unclassified samples of adversaries' malware it has discovered. CYBERCOM says the move is to improve information sharing among the cybersecurity community, but in some ways it could be seen as a signal to those who hack US systems: we may release your tools to the wider world. "This is intended to be an enduring and ongoing information sharing effort, and it is not focused on any particular adversary," Joseph R. Holstead, acting director of public affairs at CYBERCOM told Motherboard in an email.
>
> On Friday, CYBERCOM uploaded multiple files to VirusTotal, a Chronicle-owned search engine and repository for malware. Once uploaded, VirusTotal users can download the malware, see which anti-virus or cybersecurity products likely detect it, and see links to other pieces of malicious code. One of the two samples CYBERCOM distributed on Friday is marked as coming from APT28, a Russian government-linked hacking group…also known as … Fancy Bear. …CYBERCOM announced its new initiative on Monday, and uploaded its first two samples on the same day.

What do you make of CYBERCOM's decision to burn Russian tools via VirusTotal?

Project Zero:

Project Zero is part of Google. Here is a recent account, again from Motherboard:

> Ever since Project Zero was announced in 2014, these hackers have taken apart software used by millions of people—and predominantly written by other company's engineers—with a mission to "make zero-day hard." …
>
> In five years, Project Zero researchers have helped find and fix more than 1,500 vulnerabilities in some of the world's most popular software, according Project Zero's own tally. In Apple products, Beer and his colleagues have found more than 300 bugs; in Microsoft's products they found more than 500; in Adobe's Flash, they found more than 200. Project Zero has also found critical issues in CloudFlare, several antivirus apps, and chat apps such as WhatsApp and FaceTime. A Project Zero researcher was also part of the group who found the infamous Spectre and Meltdown flaws in Intel chips.
>
> These numbers show Project Zero has had a massive impact on the security of devices, operating systems, and applications used by millions of people every day.
> For Google, these disclosures give the internet giant good publicity by showing how much the company cares for the security of not just its users—but everyone else too. In assembling one of the most elite hacking teams on the planet, Google is messaging to its customers that it takes security very seriously. Along the way, Google has given itself an excuse to probe its competitors products and software, doubtlessly learning from others' security mistakes. Project Zero has been able to poke holes in the bulletproof mystique of

> the iPhone's security, which is widely believed to be the hardest consumer device to hack. In doing so, Google is able to insert itself into conversations it might not otherwise be a part of.
>
> Regardless of Project Zero's true mission, there's no doubt that the team has had a profound influence on the cybersecurity industry in the last five years.
>
> **"Without this level of technical detail in the public eye, defenders don't stand a chance."** For one, Project Zero has normalized something that years ago was more controversial: a strict 90-day deadline for companies that receive its bug reports to patch the vulnerabilities. If they don't patch in that time frame, Google drops the bugs itself. Microsoft, in particular, was not a fan of this policy at the beginning. Today, most companies that interact with Project Zero respect that 90-day deadline as an industry standard, a tidal change in the always controversial debate on the so-called "responsible disclosure"—the idea that security researchers who find vulnerabilities should first disclose them to the affected company, so that it can fix them before the bugs are exploited by hackers. According to its own tally, around 95 percent of bugs reported by Project Zero get patched within that deadline.
>
> "People looked at the way the wind was blowing and then decided that—maybe just maybe— instead of creating a fuss, creating a fix within 90 days was just easier," said Chris Evans, Project Zero's original team leader.
>
> But perhaps no accolade is more significant than how much people on the other side of Project Zero's fence, whom Evans would call the "insecurity industry," hate the Google hackers. This "insecurity industry" is made of companies like Azimuth Security and NSO Group, government contractors whose job is to find bugs and write exploits. But, instead of reporting the vulnerabilities to the companies who own the software, these companies sell them to governments who turn them into tools to hack and surveil targets.
>
> "Fuck those guys," said a researcher who works for a company that does offensive security, referring to Project Zero. "They don't make the world safer."
>
> The researcher, who spoke on condition of anonymity because they are not allowed to talk to the press, said that zero-day vulnerabilities are sometimes used to go after terrorists or dangerous criminals. So when Project Zero kills those bugs, it may be killing tools used by intelligence agencies to go after the bad guys, according to the researcher.

Consider these questions:

- Can you explain how this activity differs from VirusTotal?
- What are the pros and cons of this model?

Also check out Zero-Day Initiative here.

The MITRE ATT&CK Matrix:

Now for something different. The MITRE Corporation is a large non-profit research enterprise that houses a variety of federally funded research and development centers (FFRDCs). Put simply: it's a research institute that carries out government-funded projects. Among many other things, MITRE has an array of cybersecurity-related programs. One of these—known as the ATT&CK Matrix—has gained a remarkable following in recent years, and increasingly is used as a common touchstone

for understanding—and having a common way to talk about—the full spectrum of activities in which a malicious actor might engage. Read about it here.

Consider these questions:

- After perusing the site, can you sum up the role that the ATT&CK Matrix plays?
- Does this properly count as part of the information-sharing ecosystem?

Cyber Threat Alliance:

CTA shows us another model, one that might be described as a hybrid between the full community-good model of VirusTotal and the proprietary model associated with private sector security vendors. Here is how CTA sums up its origin and core function:

> We were founded in 2014 through an informal agreement to share intelligence among Fortinet, McAfee, Palo Alto Networks, and Symantec. They called this arrangement the Cyber Threat Alliance, but CTA had no dedicated staff nor any legal paperwork. In 2015, the companies developed a white paper on the Cryptowall Crimeware. The paper garnered a lot of attention and showed the value of collaboration among the cybersecurity community.
>
> At this point, the companies realized that they were involved in something bigger. In order to increase the impact across the ecosystem, CTA needed to scale. To achieve this, the Founding Members decided to establish CTA as an independent organization and re-launch it in February 2017 at RSA. The revamped CTA now has dedicated staff, resources, and a technology platform for sharing advanced threat data. As a result, CTA members can all share timely, actionable, contextualized, and campaign-based intelligence that can be used to improve their products and services to better protect their customers, more systematically thwart adversaries, and improve the security of the digital ecosystem.

CTA has been growing rapidly and earning a positive reputation. To understand it better, consider this account, from CTA itself, of the unusual role it can play:

> On May 23, 2018, Cisco's Talos Intelligence Group publicly exposed a new malware threat they dubbed VPNFilter. VPNFilter is a sophisticated modular malware system targeting networking equipment all over the world. This malware allowed for theft of website credentials, collection of data, injection of malicious content into network traffic as it passes through an infected device, and destruction of the infected device.
>
> Earlier in May, VPNFilter had begun a large-scale infection of devices in Ukraine and eventually targeted at least 500,000 network devices worldwide. This happened to be right around the time of the one-year anniversary of NotPetya, the Ukrainian Constitution Day, and the European Soccer Championship. The threat of a destructive attack timed with any of these events forced Talos' hand and pushed them to release all of the information they currently had on the malware to the public, in coordination with law enforcement.
>
> This is pretty typical for many cybersecurity companies. They are researching a threat and eventually decide to expose it. However, Talos did something special that has the potential to improve our cybersecurity over the long run. They informed fellow members in the Cyber Threat Alliance of the threat of VPNFilter before they released it publicly. Why would one company provide their competitors with advance notice of such a significant malware release?

Talos realized that VPNFilter had the potential to be a massive problem for internet users across the globe. If the destructive module of the malware was activated, it could cut off internet access for millions of people. They also recognized that they needed assistance in addressing the problem and informing the public of the risk. The devices that VPNFilter targeted are on the perimeter of most organizations' networks and difficult to defend, typically do not have a host-based protection system, have hundreds of publicly known vulnerabilities, and are difficult for organizations to patch. To quickly handle this threat, CTA members would need to amplify Talos' alert and spread mitigation information as quickly as possible.

Talos leveraged CTA's Algorithm & Intelligence (A&I) Committee to provide CTA members with VPNFilter cyber threat indicators and defensive measures, including VPNFilter samples and analytic findings. They also provided a briefing to members to answer questions and engage on how to mitigate the threat. CTA members were then able to quickly perform their own analysis to verify and validate Talos' work and use that to develop protections for their cybersecurity products.

When Talos released their blog at 9:00 am eastern time on May 23, CTA members already had protections in place, defending and protecting their customers as quickly as possible. This allowed our members to avoid the usual scramble and delay of trying to obtain new malware samples, performing analysis, and developing protections when another company publishes a significant report. CTA members, such as Fortinet, Symantec, Sophos, Palo Alto Networks, McAfee, Juniper, and Rapid7, published their own findings and analysis over the next few hours and days, amplifying the messaging. These blogs were collected by CTA and published in a single location for the public to find information on VPNFilter.

CTA members continued to discuss and share information on VPNFilter within A&I meetings. This included updates on victim telemetry, additional affected infrastructure, and new insights on malicious modules. Talos provided public updates on VPNFilter on June 6 and September 27, and the malware samples and analysis were shared with CTA members in advance. As before, CTA members provided additional analysis and amplification of the warnings of VPNFilter to assist with mitigation.

Impact and Assessment

Was this successful? As of the publication of this blog post, the destructive module of VPNFilter was never employed. So that's a good thing. Based on our collective visibility it appears that VPNFilter activity has been severely degraded since the release of information in May and operational coordination actions with law enforcement, intelligence organizations, and CTA and its members. Talos has seen no signs of the actor trying to reconnect with the devices that still have the Stage 1 malware, and most C2 channels for the malware have been mitigated. While it is highly unlikely that the highly capable actor behind VPNFilter has stopped their activities, it does appear that they were forced to abandon the VPNFilter framework due to these coordinated actions.

One of the other effects of Talos' early sharing and CTA's response has been an increased willingness of CTA members to share information on significant malicious cyber activity with each other before the release of details publicly. These disruption activities seek to prevent actors from succeeding in their goals and increase the costs of their malicious cyber activities. By coordinating ahead of release on significant issues, CTA members leverage their data, analysis, and cybersecurity products to expose the activity, prevent additional harm, and mitigate any of the activity's effects as early as possible. In the months since VPNFilter, Symantec (Thrip, Leafminer), Fortinet (Emotet), Sophos (SamSam), and Palo Alto Networks (Gorgon Group, KONNI, and DOGCALL) have all provided CTA members with

> early access to malicious cyber activity analysis and samples. In fact, this early release momentum continues today with Symantec's release of information and analysis on a new threat actor they call Gallmaker. CTA encourages our members to continue this type of sharing moving forward.
>
> The members of CTA seek to work together in good faith to share cyber threat intelligence to disrupt malicious actors and protect their customers. We share cyber threat indicators and defensive measures for the purposes of improving defenses against advanced cyber adversaries across member organizations and their customers, to advance the cybersecurity of critical information technology infrastructure, and to increase the security, availability, integrity, and efficiency of information systems. We recognize that not every cybersecurity provider has the same visibility into the threat, and the only way we can expand our knowledge base is to share with one another and then act for the greater good. We are all stronger when we work together. As we move forward, the early sharing demonstrated through VPNFilter provides just a glimpse of the potential impact that CTA can have to improve the overall security of the internet.

Can you articulate the value proposition CTA is describing in this story?

ISACs and ISAOs

CTA is not the only, or the first, model for cross-company collaborations involving cybersecurity information-sharing. "ISACs" and "ISAOs" are similar. The following account (from an ISAO known as the Advanced Cyber Security Center (ACSC)) provides a handy introduction:

> When it comes to sharing cybersecurity information, it can get complicated fast. Organizations have to navigate a tight workforce, conflicting laws and regulations, and prioritize protections within their organizations -- often on penny-pinching budgets. Public entities like the Federal Bureau of Investigation (FBI) and the Department of Homeland Security (DHS) strive for better engagement with corporations to improve collaboration and communication but concerns over sharing sensitive information or giving competitors insight into business operations limit collaboration.
>
> In 2015, the White House addressed those hurdles with Executive Order 13691: Promoting Private Sector Cybersecurity Information Sharing. … EO 13691 was designed to help fill the need for security information sharing beyond federal agencies and to "encourage the voluntary formation of [information sharing organizations], to establish mechanisms to continually improve the capabilities and functions of these organizations, and to better allow these organizations to partner with the Federal Government on a voluntary basis." Although the executive order provided a basis for DHS to support broadening the ecosystem, these "voluntary" sharing organizations added a layer of complexity that confused even the cybersecurity community. Information Sharing and Analysis Centers (ISACs) and Information Sharing and Analysis Organizations (ISAOs) established by the executive order differ in role and benefits but both are critically important to successful collaboration. Below, we provide a short summary of how these types of information sharing organizations operate. …
>
> **ISACs: Sector-Specific Collaboration**
>
> To address the need for better sharing of information about cyber risks and preparedness, Presidential Decision Directive 63, signed in 1998, encouraged critical

> infrastructure sectors to establish information sharing organizations. The directive established clear responsibilities at the government level for building collaborations with organizations to "serve as the mechanism for gathering, analyzing, appropriately sanitizing and disseminating private sector information" to industry and the government. Therefore, the ISACs aligned specifically with each individually designated critical infrastructure sector. ISACs provide real-time threat and attack intelligence sharing, training and curated reporting to their members, who may pay an annual fee depending on the market sector. Information sharing through ISACs is usually protected by legal agreements that protect against non-disclosure and attribution back to the reporting organization, providing some protection for their members against inappropriate disclosure. In addition, the advent of automated sharing platforms hosted by some ISACs can allow for consistent collaboration between member organizations and some may be able to respond and share actionable threat intelligence more quickly than the government. Many have also grown to reach well beyond the US, building international membership opportunities.
>
> **ISAOs: Effective-Practice Sharing Networks**
>
> While ISACs covered the emerging information sharing needs of critical infrastructure initially, the rapidly evolving cybersecurity landscape necessitated the need for Executive Order 13691 encouraging ISAO development. … [T]o encourage innovative information sharing…, "ISAOs may be organized on the basis of sector, sub-sector, region, or any other affinity, including in response to particular emerging threats or vulnerabilities." These new information sharing organizations were also envisioned to be sector-agnostic, allowing for membership to "be drawn from the public or private sectors, or consist of a combination of public and private sector organizations." Like ISACs, information sharing through ISAOs may be protected by legal agreements that protect against non-disclosure and attribution back to the reporting organization.
>
> …. ISAOs have evolved to focus beyond threat sharing and towards effective-practice sharing within a more geography-centric local community….

Task: Do some searching to see if you can identify at least three distinct ISACs or ISAOs. Come to class prepared to name them, describe their particular areas of focus, and share anything you could glean about the nature of their activities

- How, if at all, is the ISAC/ISAO model different from CTA?
- Can you articulate why some people are reluctant to see NSA involved in defense-oriented activity?
- What is the best response to such concerns?

At this point, you may have noticed that we have not yet turned our attention to some key government institutions that play a direct role in promotion of information-sharing. We've saved that for last on purpose, for the role played by DHS's Cybersecurity & Infrastructure Security Agency (CISA) goes hand-in-hand with a key statute that will require careful parsing. We'll do all of that below, in the next reading.

15. Facilitating Better Defense through Information-Sharing (II)

In this class, we conclude our survey of key information-sharing actors, and then turn our attention to an important 2015 statute designed to improve information-sharing by pruning away various disincentives.

A. More Information-Sharing Institutions

The Cybersecurity Division at the DHS Cybersecurity and Infrastructure Security Agency (CISA):

In addition to encouraging formation of ISACs and ISAOs, the federal government has taken certain institutional-design steps in order to promote information sharing. Specifically, it has charged a particular agency with this mission: the Cybersecurity and Infrastructure Security Agency (CISA, pronounced "*sis*-uh").

CISA is a component of the Department of Homeland Security (DHS). Previously known by the much less evocative moniker, the "National Protection and Programs Directorate" (NPPD), CISA has been growing rapidly both in stature and capability, and the recent name change was an important symbolic marker of that success. CISA today encompasses several distinct functional divisions, including not only its Cybersecurity Division but also its Infrastructure Security Division, the National Risk Management Center, and the Emergency Communications Division. Our primary concern is with the Cybersecurity Division.

What exactly does the Cybersecurity Division do? It has a number of distinct responsibilities. Some of these are focused on improving the security of the federal government's *own* systems. We will read about that more soon enough. Others involve cybersecurity services CISA can provide on a voluntary basis to state and local government entities and to certain private-sector entities, services that range from risk-assessment and capacity-building to penetration testing and assistance with incident response. We will table all of that for now, too, in order to stay focused on the topic of information-sharing and CISA's role in that setting.

As we noted earlier in this reading, there are several dimensions to information-sharing as an area of cybersecurity policy. Ideally, there would be an efficient flow of useful information from the federal government to the private sector, from the private sector to the government, and among various private sector entities. A variety of obstacles historically have impeded these flows, however, and thus a core challenge for CISA has been to develop useful mechanisms for overcoming those obstacles as much as possible.

How to do that? CISA in recent years operationalized these functions primarily through its 24/7 watch-and-operations center known until recently as National Cybersecurity and Communications Integration Center (NCCIC, pronounced "*N*-kick"). Since its establishment in 2009, NCCIC has served as a hub for information exchange, among other things. Those functions continue today, though recent reshufflings of internal organizational labels suggest that it may be best to eschew a focus on that precise label and to concentrate instead on simply understanding the functions CISA performs in this area.

One notable example is the *Automated Indicator Sharing* (AIS) *Program*. AIS is a free service consisting of a bi-directional information-exchange platform that facilitates real-time, automated sharing of threat indicators between CISA and voluntarily-participating private-sector entities. The basic idea is that a participating private-sector entity which identifies a threat indicator can upload signatures to CISA at machine speed via AIS, and CISA can in turn use AIS to share those signatures to all other participants in addition to using them for its own defensive purposes (note how this is akin to, say, VirusTotal). CISA can originate the dissemination as well.

Separately, CISA promotes a distinct form of information-sharing by producing (and distributing the old fashioned way) a variety of bulletins and technical advisories. Here's a current list of them, from CISA:

> Subscriptions are available to all users for the following products:
>
> - Current Activity entries provide up-to-date information about high-impact types of security activity affecting the community at large.
> - Alerts provide timely information about current security issues, vulnerabilities, and exploits.
> - Bulletins provide weekly summaries of new vulnerabilities. Patch information is provided when available.
> - Tips provide advice about common security issues for the general public.
> - Analysis Reports provide in-depth analysis on a new or evolving cyber threat.
> - Industrial Control Systems Alerts provide timely notification to critical infrastructure owners and operators concerning threats or activity with the potential to impact critical infrastructure computing networks.
> - Industrial Control Systems Advisories provide timely information about current industrial control systems (ICS) security issues, vulnerabilities, and exploits

As we will soon explore in more detail, CISA also plays an especially important information-sharing role in the special context of private sector entities that constitute "critical infrastructure." We explore that topic separately below. For now, we will move on to close out our review of key information-sharing actors by looking briefly to the National Security Agency (NSA).

NSA's New Cybersecurity Directorate:

NSA is, of course, most famous for its role as the nation's lead agency for collection and analysis of "signals intelligence" (SIGINT). In practical terms, this means that NSA is in the business of espionage targeting electronic communications, stored data, and so forth. At first blush, then, it would seem we should not be talking about NSA at all, until Unit II of this course. But NSA also has important *protective* functions related to cybersecurity. Who better to understand attackers' capabilities, after all, than our own premier electronic espionage agency? While this protective mission is largely a matter of protecting the government's *own* systems (again, a topic we will address in an upcoming class rather than here), there is an aspect of this protective mission that also concerns private-sector entities with strategic significance and the broader public. That's why we are talking about it briefly here.

The organizational home for this function, within NSA, has undergone various name and organizational-chart revisions in recent years. For a long period, it was embodied by the "Information Assurance Directorate," or IAD. After a complicated period of recent change, NSA has recently reorganized again, and now we find this function in the newly christened Cybersecurity Directorate. The general idea appears to be to maximize the effect of defensive efforts—particularly with an eye towards preventing significant intellectual-property theft by other states—by collocating a range of NSA capabilities (including collection capabilities, presumably) and creating a mission-oriented team environment for them.

The head of CD, Anne Neuberger, has explained:

> [O]ne of her directorate's new goals is to provide more actionable threat intelligence at the unclassified level so that partners, customers and private sector firms can actually reap benefits in real-time. The NSA, she said, will strive to declassify and share threat intelligence faster. . . .
>
> The Cybersecurity Directorate's early mission is to "prevent and eradicate threats" to national security and weapons systems and the defense industrial base, which face increasingly complex cyber threats from China, Russia, Iran and North Korea. . . .
>
> One of the ways the Cybersecurity Directorate aims to do that is by producing "better threat alerts with more context," Neuberger said. The directorate released such an advisory on Oct. 7, detailing to the public how multiple nation-state actors "have weaponized" certain virtual private network vulnerabilities. The advisory included a list of affected systems and recommended patches and other strategies to harden systems against intrusion.

Can you relate this to the note above regarding CYBERCOM's contributions to VirusTotal?

Having described a variety of private and public actors who do important things to promote information-sharing, we should now turn our attention to an issue that has greatly complicated that project over time: How, if at all, should Congress intervene to support such efforts via legislation intended to "prune" away disincentives to information-sharing?

B. What Do We Mean by "Pruning"?

In some situations, an entity might be willing (perhaps even eager) to pursue some particular security measure, yet may be deterred from doing so by the potential applicability of a legal constraint (criminal, civil liability, regulatory pressure, etc.). And that might be a good thing; the legal constraint presumably serves some otherwise-desirable purpose, after all, and the resulting benefits associated with that purpose might well outweigh the opportunity cost associated with discouraging uptake of the security measure in question. But then again, the balance also might cut the opposite way. If so, some "pruning"—that is, changes to the law so that it no longer will discourage the measure—might be in order.

After many years of debate, Congress in 2015 enacted a law designed to do just this in relation to the circulation of cybersecurity-relevant information. The rest of this reading examines that legislation.

C. The Cybersecurity Information Sharing Act of 2015

In 2015, Congress passed and President Obama signed into law the "Cybersecurity Information Sharing Act of 2015." Since that law has the same acronym as the newly renamed DHS agency described above, one must now be careful and clear when referring to "CISA." Here, I'll refer to it simply as the Act or the 2015 Act, in hopes of keeping things simple.

The full text of the 2015 Act can be found here, but for our purposes you should focus only on the sections quoted in the text below and the questions that follow for each.

Definitions:

As an initial matter, read these two definitions from the statute. First, Section 102(6) defines "cyber threat indicator" to mean information used to "describe or identify":

> **(A)** malicious reconnaissance, including anomalous patterns of communications that appear to be transmitted for the purpose of gathering technical information related to a cybersecurity threat or security vulnerability;
>
> **(B)** a method of defeating a security control or exploitation of a security vulnerability;
>
> **(C)** a security vulnerability, including anomalous activity that appears to indicate the existence of a security vulnerability;
>
> **(D)** a method of causing a user with legitimate access to an information system or information that is stored on, processed by, or transiting an information system to unwittingly enable the defeat of a security control or exploitation of a security vulnerability;
>
> **(E)** malicious cyber command and control;
>
> **(F)** the actual or potential harm caused by an incident, including a description of the information exfiltrated as a result of a particular cybersecurity threat;
>
> **(G)** any other attribute of a cybersecurity threat, if disclosure of such attribute is not otherwise prohibited by law; or
>
> **(H)** any combination thereof.

Second, Section 102(7) defines the term "defensive measure" as follows:

> an action, device, procedure, signature, technique, or other measure applied to an information system or information that is stored on, processed by, or transiting an information system that detects, prevents, or mitigates a known or suspected cybersecurity threat or security vulnerability … [but] does not include a measure that destroys, renders unusable, provides unauthorized access to, or substantially harms an information system or information stored on, processed by, or transiting such information system not owned by …the private entity operating the measure … or …another entity or Federal entity that is authorized to provide consent and has provided consent to that private entity for operation of such measure.

Now we will turn to the key substantive provisions in the statute. As you will see, I've reordered and grouped them in a way designed to highlight the distinct things Congress was trying to accomplish.

Two Major Goals:

It helps to think of the 2015 Act as pursuing two overarching goals. First, it sought to increase the flow of information between the government and the private sector (in both directions). Second, it sought to increase the likelihood that private sector entities would feel free to use certain defensive measures. Alas, the drafters did not organize the 2015 Act in a way that maps easily onto that two-prong framework. Accordingly, I've excerpted the relevant sections of the Act,

below, in an order that should be easier for you to follow (as compared to the mind-numbing effect of reading the sections in their original order).

Goal #1: Facilitate the Flow of Information between Government and the Private Sector

A major goal of the Act was to make it easier for the government to share certain information with the private sector, and vice-versa. Section 103 addresses the former scenario, and Section 105 addresses the latter.

a. Section 103 (When the Federal Government Shares with Others)

> **(a)** [T]he Director of National Intelligence, the Secretary of Homeland Security, the Secretary of Defense, and the Attorney General . . . shall jointly develop and issue procedures to facilitate and promote
>
> **(1)** the timely sharing of classified *cyber threat indicators* [see definition below] and *defensive measures* [see definition below] in the possession of the Federal Government with representatives of relevant Federal entities and non-Federal entities that have appropriate security clearances;
>
> **(2)** the timely sharing with relevant Federal entities and non-Federal entities of cyber threat indicators, defensive measures, and information relating to cybersecurity threats or authorized uses under this title, in the possession of the Federal Government that may be declassified and shared at an unclassified level;
>
> **(3)** the timely sharing with relevant Federal entities and non-Federal entities, or the public if appropriate, of unclassified, including controlled unclassified, cyber threat indicators and defensive measures in the possession of the Federal Government;
>
> **(4)** the timely sharing with Federal entities and non-Federal entities, if appropriate, of information relating to cybersecurity threats or authorized uses under this title, in the possession of the Federal Government about cybersecurity threats to such entities to prevent or mitigate adverse effects from such cybersecurity threats; and
>
> **(5)** the periodic sharing, through publication and targeted outreach, of cybersecurity best practices that are developed based on ongoing analyses of cyber threat indicators, defensive measures, and information relating to cybersecurity threats or authorized uses under this title, in the possession of the Federal Government, with attention to accessibility and implementation challenges faced by small business concerns.

Be able to explain how the five subparts of Section 103 accomplish different things (note how this exercise is similar to parsing the CFAA).

- What specific information-sharing obstacle do these subparts appear designed to address? That is to say: In what sense might this be described as a "pruning"?
- Note that Section 103 does not compel private-sector entities to receive whatever information the government is willing to provide. Why not take that further step?

Of course, the voluntariness of the system won't seem authentic if entities could be sued, somehow, for failing to participate.

Can you come up with an argument a plaintiff might make, along those lines?

To head off that prospect, Congress included the following language in Section 108:

> **(f)** Information Sharing Relationships.—Nothing in this title shall be construed— . . . to require the use of the capability and process within the Department of Homeland Security developed under section 105(c).
>
> **(ii)** No Liability for Non-Participation.—Nothing in this title shall be construed to subject any entity to liability for choosing not to engage in the voluntary activities authorized in this title.
>
> **(k)** Federal Preemption.— . . . This title supersedes any statute or other provision of law of a State or political subdivision of a State that restricts or otherwise expressly regulates an activity authorized under this title.

What work do these provisions do?

b. Section 105 (When Others Share with the Federal Government)

Section 105 is a very long and complicated provision. I'm going to quote from it selectively and with extensive paraphrasing, to keep this readable and focused.

First, there are two subsections (105(a) and (c)) that combine to require the federal government to take an important institutional step to facilitate the process of receiving information from others. Specifically, these subsections require the government to rapidly develop a process for receiving cyber-threat indicators and defensive measures provided by the private sector, with an emphasis on speed, automation, and distribution within the federal government once received. So far so good. But note that there was much anxiety, in some quarters, about this idea while the bill was pending. Some felt it opened the door to a privacy problem, with the government potentially obtaining personally identifiable information about customers, employees, and so forth. Accordingly, a separate subsection—105(b)—requires that the process be developed in a way that addresses such privacy and civil-liberties interests.

Note that you have seen the eventual result of Section 105, above, in the form of CISA's programs for information sharing. But there's more, so let's return to the details of the statute.

Congress anticipated that private sector entities might be reluctant to work with the government in the form of information-sharing partnerships, but it did not go so far as to mandate such cooperation. What it tried, instead, was to encourage private sector entities to participate in the system by pruning away potential disincentives. Specifically, Section 105(d) provides a host of protections, including but not limited to these three:

> **(1)** The provision of cyber threat indicators and defensive measures to the Federal Government under this title shall not constitute a waiver of any applicable privilege or protection provided by law, including trade secret protection.
>
> **(2)** [A] cyber threat indicator or defensive measure provided by a non-Federal entity to the Federal Government under this title shall be considered the commercial, financial, and proprietary information of such non-Federal entity when so designated by the originating non-Federal entity or a third party acting in accordance with the written authorization of the originating non-Federal entity.
>
> **(3)** [This one exempts the information shared by a private entity from disclosure by the government when the government receives freedom-of-information requests].

Do you think this is adequate to encourage participation? If not, what else would you be willing to try?

A separate policy challenge that emerged during the drafting of the 2015 Act concerned the possibility that the government might obtain information from the private sector for cybersecurity purposes, yet end up using the information for other purposes.

Is that a bad thing? Be able to advance at least one argument against this, and one in favor of it.

At any rate, Congress decided to address this issue. It did so in a complicated way, however. Subsection 105(d)(5) states that the information provided to the government through this mechanism may only be used for:

> **(i)** a cybersecurity purpose;
>
> **(iii)** the purpose of responding to, or otherwise preventing or mitigating, a specific threat of death, a specific threat of serious bodily harm, or a specific threat of serious economic harm . . .;
>
> **(iv)** the purpose of responding to, investigating, prosecuting, or otherwise preventing or mitigating a serious threat to a minor, including sexual exploitation and threats to physical safety; or
>
> **(v)** the purpose of preventing, investigating, disrupting, or prosecuting [certain offenses involving fraud, identity theft, espionage, censorship, and protection of trade secrets].

Is that an adequate solution (assuming a solution was needed)?

That was not the only concern Congress tried to address in subsection 105(d)(5), though. Consider this language:

> Except as provided [below], cyber threat indicators and defensive measures provided to the Federal Government under this title shall not be used by any Federal, State, tribal, or local government to regulate, including an enforcement action, the lawful activities of any [private-sector entity, except that such information can be used to inform regulatory efforts intended to support prevention/mitigation of cybersecurity threats]

Think back to our unit on regulators.

- What sort of fears might a private-sector entity have that this language tries to address?
- Does the language really address those concerns?

Next we will look at Section 104(c), titled "Authorization for Sharing or Receiving Cyber Threat Indicators or Defensive Measures":

> **(1)** In general.—Except as provided in paragraph (2) and notwithstanding any other provision of law, a non-Federal entity may, for a cybersecurity purpose and consistent with the protection of classified information, share with, or receive from, any other non-Federal entity or the Federal Government a cyber threat indicator or defensive measure.

This is followed (out of sequence) by 106(b), which provides:

> No cause of action shall lie or be maintained in any court against any private entity, and such action shall be promptly dismissed, for the sharing or receipt of a cyber threat indicator or defensive measure under section 104(c) if . . . such sharing or receipt is conducted in accordance with this title

Consider these questions:

- What legal obstacle does all this seem designed to prune?
- Does it succeed?

Goal #2: Remove Obstacles to Desirable Private-Sector Defensive Actions

Another major goal of the Act was to prune away the liability risks (real or perceived) that Congress concluded might be deterring private-sector entities from engaging in certain desirable activities (activities that are useful for defense in their own right, and that might generate important information to be shared more broadly). Most of the relevant provisions, for purposes of this goal, are found in Sections 104 and 106.

a. Making it Clear That Entities Can Monitor Their Own Systems

First, read Section 104(a):

> Notwithstanding any other provision of law, a private entity may, for cybersecurity purposes, monitor—
>
> **(A)** an information system of such private entity;
>
> **(B)** an information system of another non-Federal entity, upon the authorization and written consent of such other entity;
>
> **(C)** an information system of a Federal entity, upon the authorization and written consent of an authorized representative of the Federal entity; and
>
> **(D)** information that is stored on, processed by, or transiting an information system monitored by the private entity under this paragraph.

Absent this provision, would such activity be illegal?

Now read Section 106(a):

> No cause of action shall lie or be maintained in any court against any private entity, and such action shall be promptly dismissed, for the monitoring of an information system and information under section 104(a) that is conducted in accordance with this title.

Consider these questions:

- Are 104(a) and 106(a) duplicative?
- Note the language in 106(a) about "shall be promptly dismissed." Is that unnecessary given the prior clause stating that there shall be no cause of action?

b. Making it Clear Entities Can Do Other "Defensive" Things

Next up, consider Section 104(b), titled "Authorization for Operation of Defensive Measures." [*Tip: go back and re-read the statutory definition of "defensive measures," above, before reading 104(b)*]

> **(1)** In general.—Notwithstanding any other provision of law, a private entity may, for cybersecurity purposes, operate a defensive measure that is applied to—
>
> **(A)** an information system of such private entity in order to protect the rights or property of the private entity;
>
> **(B)** an information system of another non-Federal entity upon written consent of such entity for operation of such defensive measure to protect the rights or property of such entity; and
>
> **(C)** an information system of a Federal entity upon written consent of an authorized representative of such Federal entity for operation of such

> defensive measure to protect the rights or property of the Federal Government.
>
> **(2)** Construction.—Nothing in this subsection shall be construed—
>
> **(A)** to authorize the use of a defensive measure other than as provided in this subsection; or
>
> **(B)** to limit otherwise lawful activity.

Consider these questions:

- What activities would 104(b) encompass beyond the "monitoring" authorized by 104(a) and 106(a)?
- Imagine that a company creates a trap: a file meant to be attractive to an intruder, but includes within it malware that will, when opened, attempt to:
  - (a) communicate with the company to reveal its location,
  - (b) same as (a) but also create a backdoor through which the company can then access the system where the file is found, or
  - (c) delete all data on the system where it currently resides.
- Would any of those options be permissible under 104(b)?

c. More Pruning

Let's conclude with a look at two further pruning provisions.

Section 106(c) adds:

> Nothing in this title shall be construed . . . to create . . . (A) a duty to share a cyber threat indicator or defensive measure; or . . . a duty to warn or act based on the receipt of a cyber threat indicator or defensive measure

What liabilities does this guard against, and is it necessary?

And then we have Section 104(e):

> [I]t shall not be considered a violation of any provision of antitrust laws for 2 or more private entities to exchange or provide a cyber threat indicator or defensive measure, or assistance relating to the prevention, investigation, or mitigation of a cybersecurity threat, for cybersecurity purposes under this title.

Why might that have seemed necessary?

D. Epilogue

The 2015 Act has now been on the books for a while, and we are beginning to learn how it is working in practice. This 2018 story sounded a negative note, suggesting that relatively few companies had signed up to participate in CISA's AIS system.

> Why might that be? Give thought both to lingering obstacles and to the potential implications of the various *other* sharing systems we examined in the prior reading.

16. How the Government Protects Itself (I)

In recent classes we have surveyed the set of tools that can be used to incentivize private sector entities to adopt stronger security measures. But what about the government's *own* security practices? How might we get the government to defend *itself* better?

As with private sector entities, government entities do have internal incentives to maintain the security and functionality of their systems. No entity wants just anyone to be able to monitor its internal deliberations, for example, and plenty of entities have more particular reasons to preserve confidentiality, integrity, and accessibility (the SEC does not want people to access private information and thus enable market manipulation, for example, just as NSA does not want adversaries to be able to learn its techniques and capabilities). And yet (as opposed to the private sector), we have ample reason to believe that, if left to their own devices, many if not most government entities would not or could not invest as much in security as they probably should. And so we need tools to compel government entities, just like private entities, to do more.

Our main goal in this class is to understand which tools might be relevant to this task. Before we dive in, though, a note is in order regarding the object of our analysis.

The "government" is not a single entity, but rather a vast array of distinct agencies and enterprises (and that's true even if we concern ourselves only with the federal government, while ignoring state, county, city, tribal, and other levels of non-federal government). Most of these entities, moreover, will own and operate any number of networks, databases, and endpoints; will have their own budgetary, personnel, management, and leadership issues; and so forth. Thus, even if we limit our focus to the U.S. government (leaving aside states, counties, cities, tribal governments, territorial governments, and so forth), the number of distinct security challenges for "the government" is bewildering.

A. The 2015 OPM Hack as a Case Study

Before we dig into the various tools that might be used to push government entities to devote more effort to their own cybersecurity, it would be good to get a better grip on the problem we are trying to solve. Toward that end, then, let's start with a case study of a breach involving a government agency: the 2015 episode in which hackers (widely thought to be acting for China's Ministry for State Security) penetrated security at the federal government's Office of Personnel Management and thereby acquired a vast trove of security-clearance background-check files.

As you read about the OPM hack, bear the following in mind: Determining "how" a breach happened is a complicated matter. At one level, one can answer by explaining which vulnerabilities (technical or otherwise) were exploited to gain initial access; which tactics or methods were then used to escalate privilege, move laterally, mask presence, etc.; which

implants were installed; and how data was eventually exfiltrated, altered, or disrupted. But there's a further level that matters, too, involving how it came to be possible that this particular organization could be had in such a way. Every organization's security structure at any given point in time is a function of the organization's past choices on a number of dimensions, including: how much budget to commit to security (including personnel, hardware, software, etc.), which personnel to hire, what management structure to employ, what training to provide, what policy or legal constraints to follow, what reward (and punishment) structure to apply, and so forth. And it is a function, too, of matters that to some extent may be beyond the organization's control, such as personality dynamics and organizational culture. Both stories—the technical and the organizational—deserve attention in any case study of a particular attack. The former can yield insights for defending against similar attacks, while the latter can yield insights that might yield much broader defensive benefits.

Now, on to our case study. Please read the following excerpts from this article from Wired, which provides a handy overview of what occurred:

> The routine nature of OPM's business made the revelations of April 15, 2015, as perplexing as they were disturbing. On that morning, a security engineer named Brendan Saulsbury set out to decrypt a portion of the Secure Sockets Layer (SSL) traffic that flows across the agency's digital network. Hackers have become adept at using SSL encryption to cloak their exploits, much as online vendors use it to shield credit card numbers in transit. Since the previous December, OPM's cybersecurity staff had been peeling back SSL's camouflage to get a clearer view of the data sloshing in and out of the agency's systems.
>
> Soon after his shift started, Saulsbury noticed that his decryption efforts had exposed an odd bit of outbound traffic: a beacon-like signal pinging to a site called opmsecurity.org. But the agency owned no such domain. The OPM-related name suggested it had been created to deceive. When Saulsbury and his colleagues used a security program called Cylance V to dig a little deeper, they located the signal's source: a file called mcutil.dll, a standard component of software sold by security giant McAfee. But that didn't make sense; OPM doesn't use McAfee products. Saulsbury and the other engineers soon realized that mcutil.dll was hiding a piece of malware designed to give a hacker access to the agency's servers.
>
> . . . [E]arly that morning, a team of engineers from the US Computer Emergency Readiness Team [US-CERT], the Department of Homeland Security unit that handles digital calamities, marched into OPM's headquarters. The engineers set up a command post in a windowless storage room in the subbasement, just down the hall from where Saulsbury had discovered the hack less than 24 hours earlier. Since they couldn't trust OPM's compromised network, the visitors improvised their own by lugging in workstations and servers that they could seal behind a customized firewall.
>
> . . . One of the US-CERT team's first moves was to analyze the malware that Saulsbury had found attached to mcutil.dll. The program turned out to be one they knew well: a variant of PlugX, a remote-access tool commonly deployed by Chinese-speaking hacking units. The tool has also shown up on computers used by foes of China's government, including activists in Hong Kong and Tibet. The malware's code is always slightly tweaked between attacks so firewalls can't recognize it.
>
> . . . The hunt turned up not just malware but also the first inklings of the breach's severity. A technician from the security software company Cylance, who was supporting the effort, spotted encrypted .rar files that the attackers had neglected to delete. He knew that .rar files are used to store compressed data and are often employed by hackers to shrink files for

efficient exfiltration. In an email to Cylance CEO Stuart McClure on Sunday, April 19, the technician was blunt in his assessment of OPM's situation: "They are fucked btw," he wrote.

By Tuesday the 21st, having churned through a string of nearly sleepless days and nights, the investigators felt satisfied that they'd done their due diligence. . . . The PlugX variant they were seeking to annihilate was present on fewer than 10 OPM machines; unfortunately, some of those machines were pivotal to the entire network. "The big one was what we call the jumpbox," Mejeur says. "That's the administrative server that's used to log in to all the other servers. And it's got malware on it. That is an 'Oh feces' moment."

By controlling the jumpbox, the attackers had gained access to every nook and cranny of OPM's digital terrain. The investigators wondered whether the APT had pulled off that impressive feat with the aid of the system blueprints stolen in the breach discovered in March 2014. If that were the case, then the hackers had devoted months to laying the groundwork for this attack.

At first, the investigators left each piece of malware in place, electing only to throttle its ability to send outbound traffic; if the attackers tried to download any data, they would find themselves confined to dial-up speeds. But on April 21, Mejeur and the US-CERT team began to discuss whether it was time to boot the attackers, who would thus learn that they'd been caught. "If I miss one remote-access tool, they'll come back in through that variant, they'll reestablish access, and then they'll go dormant for six months to a year at least," says a US-CERT incident responder who participated in the OPM investigation and who agreed to speak on the condition he remain anonymous. "And then a year later, they've now put malware in a lot of different places, and you don't know what's happening because you think you already mitigated the threat."

The debate continued until the evening of Friday, April 24, when an opportunity presented itself: As part of a grid modernization program in Washington, OPM's building was scheduled to have its power cut for several hours. The team decided that, even though it would mostly be just a psychological triumph, they would dump the malware just minutes before the blackout. If the attackers were monitoring the network, they wouldn't realize their access had been cut until everything finished booting up at least 12 hours later.

By the time power was restored on the 25th, the hackers no longer had the means to roam OPM's network—or at least that's what everyone hoped. The investigators could finally turn toward piecing together what the attackers had hauled away.

. . . As the investigators laboriously sifted through interview transcripts and network logs, they created a rough timeline of the attack. The earliest incursion they could identify had been made with an OPM credential issued to a contractor from KeyPoint Government Solutions. There was no way to know how the hackers had obtained that credential, but the investigators knew that KeyPoint had announced a breach of its own in December 2014. There was a good chance that the hackers had first targeted KeyPoint in order to harvest the single credential necessary to compromise OPM.

Once established on the agency's network, they used trial and error to find the credentials necessary to seed the jumpbox with their PlugX variant. Then, during the long Fourth of July weekend in 2014, when staffing was sure to be light, the hackers began to run a series of commands meant to prepare data for exfiltration. Bundles of records were copied, moved onto drives from which they could be snatched, and chopped up into .zip or .rar files to avoid causing suspicious traffic spikes. The records that the attackers targeted were some of the most sensitive imaginable.

> The hackers had first pillaged a massive trove of background-check data. As part of its human resources mission, OPM processes over 2 million background investigations per year, involving everyone from contractors to federal judges. OPM's digital archives contain roughly 18 million copies of Standard Form 86, a 127-page questionnaire for federal security clearance that includes probing questions about an applicant's personal finances, past substance abuse, and psychiatric care. The agency also warehouses the data that is gathered on applicants for some of the government's most secretive jobs. That data can include everything from lie detector results to notes about whether an applicant engages in risky sexual behavior.
>
> The hackers next delved into the complete personnel files of 4.2 million employees, past and present. Then, just weeks before OPM booted them out, they grabbed approximately 5.6 million digital images of government employee fingerprints.
>
> . . . The Congressional hearings that take place in the wake of national calamities often have a vicious edge, and the one looking into the OPM hack was no exception. The agency's director, Katherine Archuleta, turned in a clumsy performance before the House Oversight Committee: She failed to offer a clear idea of how many people had been affected by the attack, and she seemed to duck personal responsibility by repeatedly mentioning how difficult it is to secure OPM's aging "legacy systems." The committee's members reacted with predictable scorn.
>
> . . . Damning details about OPM's porous security emerged at the hearing. The agency's own assistant inspector general for audits testified about what he characterized as a "long history of systemic failures to properly manage its IT infrastructure."
>
> The tone of the hearings struck some observers as overly brutal. The OPM brain trust received no credit for implementing the SSL decryption program that had led to the attack's discovery, nor for acting fast to quell the threat.
>
> . . .
> Archuleta resigned under pressure, and her CIO, Donna Seymour, opted for retirement days before she was to endure another round of grilling by the House committee. The two executives' departures struck fear into their peers across the federal bureaucracy.

And here are some additional details, summarizing criticism of OPM's information-security practices that were identified in a 2014 audit:

> **Unqualified InfoSec Personnel**
> The OPM's infosec system was managed by Designated Security Officers (DSO) that weren't certified IT security professionals and were performing DSO duties in addition to their full-time jobs. Despite updating to new IT security and privacy policies, the DSOs just weren't qualified to implement those policies …infosec personnel didn't report to the Chief Information Security Officer (CISO) and they didn't employ experienced infosec professionals to manage their security.
>
> . . . **Vulnerability Scanning?** . . . [A]uditors were unable to obtain tangible evidence that vulnerability scans were routinely conducted on all of OPM's servers in 2014.

No Insight, No Inventory
The report also found that OPM doesn't maintain an accurate centralized inventory of all of their servers, databases or network devices that reside within the network. Without insight into this basic data, it's pretty hard to ensure OPM data is secured.

Untimely Patch Management
Although the OCIO applies operating system patches on all devices within OPM's network weekly, and uses a third-party patching software management program to update the software, they still found, via scans, that numerous servers were not patched on a timely basis.

. . . **Poorly Configured SIEM**
Although the OPM owns a security information and event management (SIEM) tool that can detect, analyze and correlate security incidents over time, it wasn't configured to receive data from 20 percent of major OPM information systems.
The report also stated that OPM systems were over-reporting log and event data, resulting in too much data for their security analysts to review. Plus, there was a high volume of false-positives that created a backlog and delay in identifying real incidents.

No PIV Authentication
While over 95 percent of OPM workstations require PIV (Personal Identity Verification) authentication to access the OPM network, none of the agency's 47 major apps require PIV authentication. PIV is a type of government-issued smart card.

No Two-Factor Authentication
According to the NYTimes.com, in an interview, OPM's Chief Information Officer (CIO) Donna Seymour said that installing two-factor authentication (multi-factor authentication) in the government's 'antiquated environment' was difficult and very time-consuming.

. . . Ultimately, the security report found that OPM had not undergone an adequate security controls test in more than eight years - which is more than enough time for new vulnerabilities to pop up and for breaches to go unnoticed. Don't wait that long to take a look at your own company's security controls in order to avoid a preventable breach.

Consider the following questions:

Focus first on the technical story:
- How did the intruders gain access to OPM's system in the first place?
- How were they able to move laterally and extract data without detection?

Now focus on the organizational story:
- Can you identify (or at least imagine) a set of factors that help explain why OPM did not have better security?
- For each factor you identified, what exactly would a "fix" look like?
- What factors might make it hard in practice to adopt each such fix?

As long as we have this case study cued up, let's use it for a quick review:
- Can you categorize what China apparently did here, using the typology of government actions we reviewed earlier in the course?
- In light of that categorization, and drawing on our prior discussions of US-China relations, how should the U.S. government have responded to the OPM hack? Think of three specific potential actions, and list pros and cons for each.

B. Tools to Generate Better Government Security

Plainly, the government's own cybersecurity efforts sometimes fall short. And so the question arises: What tools can we bring to bear to get the government to try harder?

That's the same question we have been exploring throughout Unit I.B., with reference to private-sector entities as the victims. We therefore might start by considering whether the tools used to push the private sector to do better might also help the government to push itself to do better.

1. Tools that Work Poorly in this Setting

As we saw previously, one way to encourage defensive improvements is to increase the extent to which entities perceive that they are exposed to the risk of enforcement actions by regulators such as the FTC. Another is leverage from insurance companies. Neither really work where it is the government itself we are trying to impact.

> Can you explain why that might be?

Government entities *do* face the possibility of suits brought by private plaintiffs, at least in theory. But such suits have a bad track record.

The first problem is "sovereign immunity." Unlike a private entity, federal and state government entities cannot be hauled into their own courts involuntarily; as sovereigns, they can only be sued if they have consented.

Does that ever happen? Yes, in fact waiver of sovereign immunity via statute is common. Both federal and state laws are full of examples of statutes expressly waiving immunity as to certain types of claims. The question for us is: Do any of them apply in a setting where the government entity had poor cybersecurity? There are some that *might*, but in practice they've not yet proven to have much bite in the cybersecurity setting.

The most well-known statute of this kind at the federal level is the Federal Tort Claims Act. It waives immunity where a person suffers personal injury, property damage, or death as a result of wrongful or negligent conduct by a government official. Most states (including Texas) have something similar on the books. It is difficult to use these laws to sue successfully for damages relating to a data breach, however, in light of the injury requirement.

A second option is the federal Privacy Act, which attempts to protect the confidentiality of personally identifiable information held by the U.S. government. In relevant part, 5 U.S.C. 552a(e)(10) states that the government must:

> establish appropriate administrative, technical, and physical safeguards to insure the security and confidentiality of records and to protect against any anticipated threats or hazards to their security or integrity which could result in substantial harm, embarrassment, inconvenience, or unfairness to any individual on whom information is maintained . . . .

The Privacy Act is a key part of the plaintiffs' claims in currently pending litigation known as *In re: U.S. Office of Personnel Management Data Security Breach Litigation*, which is a would-be class-action stemming from the OPM breach. The plaintiffs argue that OPM's poor security measures violate the Privacy Act (among other things). Here, too, the question of damages looms large. The Supreme Court in 2012's *FAA v. Cooper* had held, after all, that "emotional damages" are not

cognizable under the Privacy Act. And with the OPM hack attributed widely to China as an act of espionage rather than an effort to raise money via identity theft, it was not surprising when the trial court initially dismissed the lawsuit on the ground that the claims of damage were speculative, and thus that the plaintiffs lacked standing:

> Rejecting plaintiffs' argument that they faced a heightened risk of identity theft due to the breaches, the court held that the facts alleged failed to plausibly support the conclusion that this risk of future injury was either substantial or clearly impending. The district court ultimately concluded that only those plaintiffs who specifically identified out-of-pocket losses stemming from the actual misuse of their data had suffered an injury in fact sufficient for standing purposes. But even those plaintiffs lacked standing, the district court concluded, because they failed to allege facts demonstrating that the misuse of their information was traceable to the OPM breaches in particular.

But then, in the summer of 2019, a divided D.C. Circuit Court of Appeals partially reversed this determination. The majority explained:

> [T]he district court should not have relied even in part on its own surmise that the Chinese government perpetrated these attacks. Absent any factual allegations regarding the identity of the cyberattackers, the district court was not free to conduct its own extra-record research and then draw inferences from that research in OPM's and KeyPoint's favor.
>
> . . . Beyond that, although a cyberattack on a government system might well be motivated by a purpose other than identity theft, given the type of information stolen in the OPM breaches and Arnold Plaintiffs' allegations regarding the subsequent misuse of that information, it is just as plausible to infer that identity theft is at least one of the hackers' goals, even if those hackers are indeed affiliated with a foreign government. Our dissenting colleague takes a different tack, suggesting that because this case involves government databases, "espionage is an 'obvious alternative explanation'" for the attacks. . . . We disagree as to just how obvious an explanation this is based on the facts alleged in the complaint. Furthermore, given that espionage and identity theft are not mutually exclusive, the likely existence of an espionage-related motive hardly renders implausible Arnold Plaintiffs' claim that they face a substantial future risk of identity theft and financial fraud as a result of the breaches.
>
> . . . Because Arnold Plaintiffs adequately allege a substantial risk of future identity theft, any expenses they have reasonably incurred to mitigate that risk likewise qualify as injury in fact.

Are you persuaded? Why, or why not?

Intriguingly, some of the plaintiffs in the OPM litigation argued that OPM's poor effort to protect their information violated not just the Privacy Act but also the *Constitution* itself. Specifically, they argued that there is an (unwritten) constitutional right to "informational privacy" that can be derived from the Due Process Clause of the Fifth Amendment (the text of which prohibits the deprivation of life, liberty or property without due process of law), and that this asserted right perhaps could be read to impose an obligation of reasonable efforts to protect personally identifiable information once in the government's hands.

Is there such a right in the first place? The Supreme Court has largely dodged the question. In its 2009 *NASA v. Nelson* decision, the majority assumed without deciding that some such a right might exist, citing two earlier cases that it depicted as having made a similar move (Justices Scalia and Thomas dissented, calling for clear recognition that no such right has as yet been made part of the Constitution, however desirable such a right might be as a matter of policy). But some circuit courts, such as the Ninth Circuit, have been willing to go further, and in the same D.C. Circuit decision in the OPM case discussed above, the majority was willing to assume, at least for now, that such a claim exists.

Consider these questions:

- As a normative matter, what are the best arguments for and against permitting persons to sue for damages in circumstances like the OPM breach?
- If such a right should exist, should it take the form of a statute or of a constitutional right?
- Let's assume the government can be sued in these cases. Do you think that anticipation of liability risk for the government as a whole will impact agency decision-making in a manner similar to the way that litigation risk may impact the decision-making of private entities? Put yourself in the position of an agency head making budget plans. What are your incentives, exactly?

2. Tools that Work Well in this Setting

Whereas regulatory enforcement, insurance leverage, and liability risk all work better for influencing the private sector than for influencing the government itself, the reverse is true with respect to another tool: direct mandates to do (or not do) certain things.

Why would it be easier for the government to impose direct mandates on itself, as compared to imposing them on the private sector?

Our task now is to understand the way that the federal government, over time, has developed a complex system for imposing security mandates on itself.

Let's start (somewhat arbitrarily) in 1996, during the Clinton administration. A statute passed that year tasked the Secretary of Commerce with establishing information-security standards that *most* of the rest of the government would henceforth have to follow (the obligation would not extend to defense and intelligence agencies, however; they remained free to develop and enforce their own rules in this respect). The statute specified that the Secretary should base his or her directives on the standards and guidelines developed by the Commerce Department's National Institute of Standards and Technology (better known as "NIST"), which is a deeply respected technical organization.

Fast-forward six years, to the early years of the George W. Bush administration. In 2002, at a time before the Department of Homeland Security existed, Congress passed the Federal Information Systems Management Act of 2002 ("FISMA," pronounced "fiz-muh"). FISMA did three significant things, for our purposes. First, it took the standard-setting role away from Commerce and gave it to the White House's Office of Management and Budget (OMB) (though in fulfilling that role OMB still would remain reliant on the guidance provided by NIST, which remained part of the Commerce Department). Second, FISMA charged OMB with a critical further task: to monitor, on

at least an annual basis, whether other agencies actually complied with its standards. And third, FISMA called for creation of an "information security incident center" that would both provide expert advice (including in the face of an unfolding emergency) and function as a threat-intelligence hub (collecting and analyzing information). This eventually resulted in creation of US-CERT, which you read about previously in relation to the Equifax and OPM case studies.

Six years later, with DHS now in existence, President Bush issued National Security Presidential Directive 54/Homeland Security Presidential Directive 23. That 2008 Directive called for DHS to play the lead role in protecting federal networks (other than military and intelligence systems). This meant something much more concrete and direct than simply generating rules for other agencies to follow and then monitoring to see whether they complied. The order made US-CERT part of DHS (in particular, part of the directorate that eventually became CISA), and directed DHS to act through US-CERT to monitor and protect all "external access points" associated with federal government systems, and to provide intrusion detection, incident analysis, and other capabilities. That is to say, the order called upon DHS to take on operational responsibility for certain security functions on behalf of all other non-military, non-intelligence agencies. DHS responded to that assignment by, among other things, developing and deploying an intrusion-detection system that came to be known as "Einstein." Here is a DHS account of Einstein versions 2 and 3 from a few years ago:

> DHS is deploying, as part of its EINSTEIN 2 activities, signature-based sensors capable of inspecting Internet traffic entering Federal systems for unauthorized accesses and malicious content. The EINSTEIN 2 capability enables analysis of network flow information to identify potential malicious activity while conducting automatic full packet inspection of traffic entering or exiting U.S. Government networks for malicious activity using signature-based intrusion detection technology. . . . EINSTEIN 2 is capable of alerting US-CERT in real time to the presence of malicious or potentially harmful activity in federal network traffic and provides correlation and visualization of the derived data. [Meanwhile, DHS is developing a new system], called EINSTEIN 3, [that] will draw on commercial technology and specialized government technology to conduct real-time full packet inspection and threat-based decision-making on network traffic entering or leaving these Executive Branch networks. . . . The EINSTEIN 3 system will also support enhanced information sharing by US-CERT with Federal Departments and Agencies by giving DHS the ability to automate alerting of detected network intrusion attempts and, when deemed necessary by DHS, to send alerts that do not contain the content of communications to the National Security Agency (NSA) so that DHS efforts may be supported by NSA exercising its lawfully authorized missions. . . . DHS will be able to adapt threat signatures determined by NSA in the course of its foreign intelligence and DoD information assurance missions for use in the EINSTEIN 3 system in support of DHS's federal system security mission. Information sharing on cyber intrusions will be conducted in accordance with the laws and oversight for activities related to homeland security, intelligence, and defense in order to protect the privacy and rights of U.S. citizens.

Can you sum up what the Einstein system actually does, and for whom?

With DHS thus emerging as the lead agency for most aspects of the federal government's own cybersecurity, a question arose regarding whether it still made sense for OMB to have the oversight role assigned to it by the 2002 FISMA (that is, the role of ensuring that other agencies were complying with OMB's NIST-based standards). In 2010, the Obama administration concluded that the answer was no, and notwithstanding FISMA, proceeded to direct OMB to delegate its oversight role to DHS. Then, in 2014, Congress ratified that division of labor in "FISMA 2014." In addition, FISMA 2014 also enhanced DHS's authority in a crucial respect: from now on, DHS

(through what is now CISA) would have the ability not only to take note of (and complain about) an agency's failure to comply with OMB's NIST-based rules, but also the ability to issue "binding operational directives" requiring agencies to take some particular security-related step.

CISA issued its first BOD in May 2015. For a list of the BODs that CISA has issued, see here. One that drew a great deal of attention occurred in September 2017. In BOD 17-01, CISA compelled government agencies to take steps to "remove and discontinue present and future use of all Kaspersky-branded products." (Note: if you are interested, you can read the memorandum explaining this decision here). Not long after CISA acted, Congress effectively adopted this BOD in the form of a statutory ban.

Consider these questions:

- What are the pros and cons of BODs versus statutes, and do we need both options?
- Recall that the OPM hack took place in 2015. Does the overall description of the government's self-regulatory approach over time, set forth above, impact your view regarding lessons to be learned from the OPM fiasco?

Before we move on, a quick note about the current status of US-CERT is in order. You may recall from our initial introduction to CISA that the functions of NCCIC remain in place even as internal organizational realignments call into question the continuing relevance of that specific label. Well, suffice to say that much the same can be said about US-CERT: the functions of US-CERT remain as before, and the name itself is still in use in certain respects, but it is no longer clear that the old label is the best way to capture how CISA itself refers to these functions. For those of us on the inside, therefore, the same advice obtains: concentrate on learning the functionality, and be mindful that different people may be using different organizational labels to refer to them.

And now, back to our chronology of the government's evolving approach to self-regulation. The most recent major change came in May 2017, when President Trump issued Executive Order 13,800 (May 11, 2017), "Strengthening the Cybersecurity of Federal Networks and Critical Infrastructure." For now, we are concerned only with Section 1 of that order. In relevant part, it states:

> **(c)** Risk Management
>
> **(i)** Agency heads will be held accountable by the President for implementing risk management measures commensurate with the risk and magnitude of the harm that would result from unauthorized access, use, disclosure, disruption, modification, or destruction of IT and data. They will also be held accountable by the President for ensuring that cybersecurity risk management processes are aligned with strategic, operational, and budgetary planning processes.
>
> **(ii)** Effective immediately, each agency head shall use The Framework for Improving Critical Infrastructure Cybersecurity (the Framework) developed by the National Institute of Standards and Technology, or any successor document, to manage the agency's cybersecurity risk. Each agency head shall provide a risk management report to the Secretary of Homeland Security and the Director of the Office of Management and Budget (OMB) within 90 days of the date of this order. The risk management report shall:

> **(A)** document the risk mitigation and acceptance choices made by each agency head as of the date of this order, including:
>
> **(1)** the strategic, operational, and budgetary considerations that informed those choices; and
>
> **(2)** any accepted risk, including from unmitigated vulnerabilities; and
>
> **(B)** describe the agency's action plan to implement the Framework.
>
> **(iii)** The Secretary of Homeland Security and the Director of OMB…shall jointly assess each agency's risk management report to determine whether the risk mitigation and acceptance choices set forth in the reports are appropriate and sufficient to manage the cybersecurity risk to the executive branch enterprise in the aggregate (the determination).

In what precise way does this approach increase pressure on agencies?

Before we close this discussion, consider one further "tool" with which an entity might be incentivized to put more effort into defending itself: a version of the "name-and-shame" concept.

Does it surprise you to see name-and-shame mentioned here? We talked about this concept a great deal in Unit I.A, where we focused on tools for imposing costs on attackers. But how can it apply usefully in Unit I.B, where our goal instead is to encourage or otherwise enable better defense by potential victims?

The answer is that useful incentives may be created when relevant decisionmakers within an organization understand that they may in some sense be blamed personally in the event of a security breach. If the decisionmaker anticipates such individual accountability and believes negative consequences will attach to it, in other words, the person may prove willing to give greater weight to security considerations.

What might the negative consequences be? There are many possibilities: public shaming in the form of criticism in the media or even in-person before a Congressional committee; private shaming within the organization; blunted career prospects; demotion; even termination from employment. Consider the fate of key personnel from the OPM story for illustrations (and note that much the same can be true in the private sector).

From this perspective, can you mount an argument that the notoriety of the OPM fiasco was good for government cybersecurity going forward?

17. How the Government Protects Itself (II)

In our last class, we surveyed the tools that do and do not work well when it comes to pushing the government to protect itself better. Along the way, we noted at several points that the roles

played by OMB and DHS/CISA do *not* apply to systems belonging to the military or to the Intelligence Community. Our task now, accordingly, is to understand at least a bit about how those critical parts of the federal government go about defending themselves.

A. Cybersecurity for the Military

The efforts of the Department of Defense ("DoD") to defend itself can be broken down into two categories: First, direct defense of its own networks and other systems that might constitute "blue space." Second, indirect efforts to boost the defense of private-sector actors that play an important role in relation to national defense.

1. Securing the DoDIN

Not surprisingly, DoD is not subject to the OMB/CISA system of cybersecurity oversight that we discussed in the last reading. Instead, DoD itself shoulders the responsibility for ensuring that its own vast web of information systems and communication networks are defended appropriately.

That vast web is known, collectively, as the Department of Defense Information Network, or "DoDIN" (pronounced "**doh**-din"). One commentator memorably described DoDIN as "really not a single network, but a quasi-feudal patchwork of often incompatible local networks. It's the Holy Roman Empire of cyberspace."

How does DoD organize to ensure the DoDIN is defended appropriately?

Start with the fact that the constituent elements of the DoDIN each belong to some constituent element within DoD: a combatant command (for example, Central Command (CENTCOM)), a service branch (for example, the Navy), a combat-support agency (for example, the Defense Information Support Agency), and so forth. As Joint Publication 13-2 explains, each such entity is authorized and has some capacity of its own "to take . . . defensive actions . . . and to restore the system to a secure configuration." Not all threats can be dealt with sufficiently by such "local" resources, however, and in any event, many threats are not confined to a single entity's systems. This, plus the general need for higher levels of planning, coordination, and oversight, eventually gave rise to the creation of a headquarters element focused on DoDIN defense: "JFHQ-DoDIN," which stands for Joint Forces Headquarters-DoDIN.

JFHQ-DoDIN is a "subordinate headquarters" within U.S. CYBERCOM (an organization we have discussed some already, and will discuss more once we reach Unit II). It can assign "cyber protection teams" as needed to assist in defense of the DoDIN, and it also can take charge or coordinate as needed where multiple elements of the system are involved, or where there may be spillover effects of other kinds.

In addition to this operational function, JFHQ-DoDIN also has an oversight/management function analogous to that performed by OMB/CISA. On that dimension, JFHQ-DoDIN oversees the owner/operators of the many component parts of the DoDIN as they engage in prospective risk-assessment; aggregates and assesses the results from those efforts; and issues directives for change where needed.

2. Beyond the DoDIN: Defending "Blue Space"

DoD uses the phrase "blue space" to describe cyber operational environments in which DoD is a known and authorized presence. This includes, obviously, the DoDIN itself. But there are contexts

in which DoD may become a known and authorized defender. Collectively, the combination of the DoDIN with these others constitute "blue space" environments.

"Blue space" is in contrast to operations that might occur on an adversary's system ("red space," which includes any system the adversary has the ability to control to the exclusion of others), and operations that might occur on the systems of third parties ("gray space," which includes any system not constituting blue space or red space).

> Can you think of reasons why these distinctions matter?

Of course, the idea of DoD being asked to defend non-DoDIN systems raises a question about when and how such requests occur. In contexts involving partner nations overseas, it typically is as a government-to-government agreement. CYBERCOM personnel, for example, have assisted countries including Montenegro, Ukraine, and North Macedonia via "threat hunting" missions intended to identify malware targeting critical systems there. Such missions not only benefit those allies, but also directly benefit the United States (and other allies) because they generate valuable threat intelligence (IOCs, insights into tactics, etc.). Indeed, insofar as Russia in particular tests new tools and tactics on countries like North Macedonia before using them against U.S. targets, this approach can function as an important early warning system.

When the context for non-DoDIN blue-space operations instead is domestic, things of course are more complicated given the robust traditions, legal constraints, and political sensitivities that combine to limit the role that the U.S. military normally can play in domestic life. That said, there are contexts in which DoD personnel or equipment are used to assist civil authorities during an exigency, particularly where the military has unique capabilities. DoD refers to this category of activity, in general, as "Defense Support of Civil Authorities" (DSCA). It is the sort of thing that comes up without much controversy during hurricanes or other natural disasters, but it can and does extend to other contexts. These days, one such context is cyberspace.

The following passage, from a recent Government Accountability Office report, provides an easy introduction to two statutory authorities that might result in CYBERCOM providing support to civil authorities:

> Under the Robert T. Stafford Disaster Relief and Emergency Assistance Act (**the Stafford Act**), when state capabilities and resources are overwhelmed and the President of the United States declares an emergency or disaster, the governor of an affected state can request assistance from the federal government for major disasters or emergencies. The Stafford Act aims to provide a means of assistance by the federal government to state and local governments in responding to a presidentially declared major disaster or emergency. A governor's request for the President to declare a major disaster or emergency is required to be based on a finding that the situation is of such severity and magnitude that effective response is beyond the capabilities of the state and the affected local governments and that federal assistance is necessary.
>
> Additionally, under **the Economy Act**, a federal agency may request the support of another federal agency, including DOD, without a presidential declaration of a major disaster or an emergency. This act permits one federal agency to request goods and services from another federal agency provided that, among other things, the service is available and cannot be obtained more cheaply or conveniently by contract.

Consider these questions:

- Can you explain why it is easier to rely on the Economy Act than the Stafford Act?
- Can you explain how each might be the vehicle allowing CYBERCOM to extend its defensive capabilities to the benefit of CISA, a state, or another entity?

Note: Any conversation on this topic inevitably comes around to some version of the question: Doesn't the Posse Comitatus Act (PCA) bar the military from having any involvement in domestic affairs? That is a common, but incorrect, understanding of the PCA. It provides:

> Whoever, except in cases and under circumstances expressly authorized by the Constitution or Act of Congress, willfully uses any part of the Army or the Air Force as a posse comitatus or otherwise to execute the laws shall be fined under this title or imprisoned not more than two years, or both.

Can you explain why, in light of that language, CYBERCOM at least sometimes can provide support to CISA, etc.? Try to identify two distinct reasons.

B. Protecting the Defense Industrial Base

Earlier in this Unit, we took note of the extent to which insurers and other contract counterparties can drive organizations to improve their security posture simply by insisting on various terms as conditions for the contract. At that time, we noted that there are contexts in which the federal government in particular can have a broad impact of this kind (just think of how the government has leverage through size when it comes to bargaining down prices in the context of Medicaid and Medicare). One very good example of this arises with respect to DoD and its use of contract leverage to drive improvements to security among businesses that are part of the "Defense Industrial Base" ("DIB"), as well as businesses elsewhere in DoD's many supply chains.

Here is the (over)simplified version of a painfully complex topic: DoD has a set of regulations that govern its "acquisition" processes, known as the Defense Federal Acquisition Regulation Supplement (or simply "DFARS"). In 2016, a rule was added to DFARS that required would-be contractors to be in compliance with a set of procedures that NIST had developed with an eye towards improving cybersecurity amongst critical infrastructure owners (more on them in our next reading). Basically, this meant promising to adopt the practices NIST recommended, and to document compliance. The new rule also required these companies to query their own supply chains, moreover, to find out whether they complied with the NIST framework (and if not, why not). Finally, the rule also required contractors to give DoD notice within 72 hours in the event of a cyber incident.

Here's the catch: there was no mechanism for actually vetting these representations before the contract was let, and only limited resources available for post-agreement auditing (to be performed by the Defense Contract Management Agency (DCMA)). In short, it was really just a trust model of compliance.

Not surprisingly, an investigative report from MITRE recently concluded that compliance in fact was quite limited. Citing this, DoD began exploring a more aggressive approach, involving something called the "Cybersecurity Maturity Model Certification" (CMMC). The idea here is (1)

to develop a sliding-scale of security obligations that a would-be contactor will have to meet (with the scale demanding more from the entity in accordance with the sensitivity of the data the entity will have), and—here's the real novelty—(2) each such business must be certified for compliance before the contract can be awarded.

Consider these questions:

- The certification requirement presents a scalability issue. Can you explain how so?
- It seems likely that third-party businesses will fulfill much of the certification role. What are the pros and cons of this, as compared to expecting DCMA or some other government body to handle all the work?

Note: If you are interested and wish to go deeper, read this for a review of some other key questions that remain open at this time.

C. Cybersecurity for the Intelligence Community

What about protection of the information systems that carry the nation's classified information and other aspects of the intelligence enterprise, apart from the DoDIN? A key term-of-art for such systems is "National Security Systems." As noted above, the OMB/CISA role does not apply to such systems, but they are not the exclusive province of the military either. Who defends them?

The answer is complicated, for various agencies within the Intelligence Community, of course, have their own in-house capabilities with respect to the defensive aspects of cybersecurity. But for our purposes, the main thing to understand is that the NSA plays the most significant operational role in protecting these networks.

Related to this is the Committee on National Security Systems. CSSS is an interagency body with representation from 21 separate government agencies, tasked with setting national policy—and issuing corresponding guidance and procedures—with respect to protection of National Security Systems. NSA's Director is the designated "National Manager" for this function, as one might expect. Note that CSSS can and does issue directives with which agencies must comply, much as CISA can for civilian systems and JFHQ-DoDIN can for military systems.

C. Mitigation and Resilience: Consequence Management

So far in this course we have examined both (i) tools for imposing costs on attackers and (ii) tools for driving potential victims to improve their security. No matter how effective our policies may be on both dimensions, though, some attacks will get through. What then? It's time to talk about "incident response" in the event of a breach.

There are many lenses one can bring to bear on this topic. First and foremost, of course, are the technical challenges that arise once one realizes that some sort of breach has occurred. We caught a glimpse of this earlier when we read about the OPM breach. As important as this topic is, however, it is beyond the scope of our course, so having noted it we will move along so we can focus on legal, policy, and institutional questions.

The next lens to note concerns the mix of competing non-technical considerations that plague decision-making for the leaders of organizations that have suffered a breach. Again, we actually have touched upon this already, in the course of our discussion of liability for violation of breach

notification laws. In that reading, we noted that there is a serious tension among several interests that an entity may have in that situation. Complicating matters, some of these interests are aligned with the public's interests, but some are not. The interests include:

(a) providing potentially impacted persons (those whose data may have been exposed) with rapid notice so that they can take steps to mitigate harm; and

(b) preserving secrecy in order to facilitate

> (i) investigation into the source, nature, and extent of the breach (including both the entity's own investigation as well as any law enforcement or intelligence investigative activity);
>
> (ii) mitigation and remediation efforts (*i.e.*, efforts to boot out and keep out the attacker); and
>
> (iii) avoiding reputational harm and other such costs that might arise.

Having reminded you of those tensions, we will not further revisit that topic either. Instead, we will now turn to a third lens associated with incident response: the question of how the *government* has organized itself to take action with respect to breaches of special, national significance.

Bear in mind that breaches come in all shapes and sizes, and in most instances, the consequence-management challenge is a matter of concern primarily to the victim entity itself (as well as to those persons whose data may have been exposed). Sometimes, though, a breach has wider significance—perhaps even calling for involvement by the U.S. government using one or more of the cost-imposition tools we discussed in Unit I.A; providing defensive, investigative, or mitigation support; or both. Since it is not always self-evident when the U.S. government should become involved, though, let alone which tools it should employ and which of its many institutions should take action, it stands to reason that we should have some form of coordination mechanism.

Our primary aim in this one-class subunit is to consider which situations warrant such involvement, what form such involvement might take, and how the government has organized itself to answer those questions when particular cases arise.

18. Critical infrastructure (I)

Most breaches do not implicate the national interest, at least not when considered in isolation (the net effect of breaches impacting intellectual property is a different story, of course). Put another way, run-of-the-mill incidents—particularly those involving the private sector—may warrant a law-enforcement response of some kind, but normally they do not warrant consideration of the larger question of whether and how to marshal and deploy the various instruments of national power.

Yet some scenarios *do* warrant exactly that. Our task: understanding which contexts "count" in this particular way; which officers or institutions get to decide when such a context is present; and what follows from such a determination.

A. An Introduction to Critical Infrastructure

A good place to begin with this topic is to establish familiarity with the idea of "critical infrastructure" (that is, "CI"). That formerly obscure label is now commonplace. But what does it really mean?

At a high level of generality, "CI" captures that idea that our daily lives and the economy depend to no small extent on certain particularly important systems, services, and structures. As DHS has put it:

> [C]ritical infrastructure provides the essential services that underpin American society and serve as the backbone of our nation's economy, security, and health. We know it as the power we use in our homes, the water we drink, the transportation that moves us, the stores we shop in, and the communication systems we rely on to stay in touch with friends and family.

Consider these questions:

- Are the generalized examples mentioned in that second sentence all equally "critical"? If not, rank them from most vital to least as you see it.
- What policy implications follow from the idea that some CI is more critical than others?

Now let's get more specific. A federal statute (42 U.S.C. 5195c(e)) defines "critical infrastructure" to mean:

> [S]ystems and assets, whether physical or virtual, so vital to the United States that the incapacity or destruction of such systems and assets would have a debilitating impact on security, national economic security, national public health or safety, or any combination of those matters.

In 2013, moreover, President Obama via Presidential Policy Directive 21 ("PPD-21") split the general category of CI into 16 distinct "sectors." Here is the list as set forth on CISA's page on this topic:

> 1. Chemical Sector
> The Department of Homeland Security is designated as the Sector-Specific Agency for the Chemical Sector.
>
> 2. Commercial Facilities Sector
> The Department of Homeland Security is designated as the Sector-Specific Agency for the Commercial Facilities Sector, which includes a diverse range of sites that draw large crowds of people for shopping, business, entertainment, or lodging.
>
> 3. Communications Sector
> The Communications Sector is an integral component of the U.S. economy, underlying the operations of all businesses, public safety organizations, and government. The Department of Homeland Security is the Sector-Specific Agency for the Communications Sector.
>
> 4. Critical Manufacturing Sector
> The Department of Homeland Security is designated as the Sector-Specific Agency for the Critical Manufacturing Sector.
>
> 5. Dams Sector

The Department of Homeland Security is designated as the Sector-Specific Agency for the Dams Sector. The Dams Sector comprises dam projects, navigation locks, levees, hurricane barriers, mine tailings impoundments, and other similar water retention and/or control facilities.

6. Defense Industrial Base Sector
The U.S. Department of Defense is the Sector-Specific Agency for the Defense Industrial Base Sector. The Defense Industrial Base Sector enables research, development, design, production, delivery, and maintenance of military weapons systems, subsystems, and components or parts to meet U.S. military requirements.

7. Emergency Services Sector
The Department of Homeland Security is designated as the Sector-Specific Agency for the Emergency Services Sector. The sector provides a wide range of prevention, preparedness, response, and recovery services during both day-to-day operations and incident response.

8. Energy Sector
The U.S. energy infrastructure fuels the economy of the 21st century. The Department of Energy is the Sector-Specific Agency for the Energy Sector.

9. Financial Services Sector
The Department of the Treasury is designated as the Sector-Specific Agency for the Financial Services Sector.

10. Food and Agriculture Sector
The Department of Agriculture and the Department of Health and Human Services are designated as the co-Sector-Specific Agencies for the Food and Agriculture Sector.

11. Government Facilities Sector
The Department of Homeland Security and the General Services Administration are designated as the Co-Sector-Specific Agencies for the Government Facilities Sector.

12, Healthcare and Public Health Sector
The Department of Health and Human Services is designated as the Sector-Specific Agency for the Healthcare and Public Health Sector.

13. Information Technology Sector
The Department of Homeland Security is designated as the Sector-Specific Agency for the Information Technology Sector.

14. Nuclear Reactors, Materials, and Waste Sector
The Department of Homeland Security is designated as the Sector-Specific Agency for the Nuclear Reactors, Materials, and Waste Sector.

15. Transportation Systems Sector
The Department of Homeland Security and the Department of Transportation are designated as the Co-Sector-Specific Agencies for the Transportation Systems Sector.

16. Water and Wastewater Systems Sector
The Environmental Protection Agency is designated as the Sector-Specific Agency for the Water and Wastewater Systems Sector.

Note the references above to "sector-specific agencies," as we will return to that in a moment.

> Would you add anything to this list? What do you think does *not* belong?

PPD-21 also directed the Secretary of Homeland Security to update this list periodically. Normally, such moves draw little attention. But just prior to the inauguration of President Trump in January 2017, DHS Secretary Jeh Johnson added election systems to the CI list for the first time (adding this as a subsector in the government-facilities category, alongside existing subsectors for "education facilities" and "national monuments and icons"). Secretary Johnson explained:

> I have determined that election infrastructure in this country should be designated as a subsector of the existing Government Facilities critical infrastructure sector. Given the vital role elections play in this country, it is clear that certain systems and assets of election infrastructure meet the definition of critical infrastructure, in fact and in law. . . .
>
> By "election infrastructure," we mean storage facilities, polling places, and centralized vote tabulations locations used to support the election process, and information and communications technology to include voter registration databases, voting machines, and other systems to manage the election process and report and display results on behalf of state and local governments.
>
> Prior to reaching this determination, my staff and I consulted many state and local election officials; I am aware that many of them are opposed to this designation. It is important to stress what this designation does and does not mean. This designation does not mean a federal takeover, regulation, oversight or intrusion concerning elections in this country. This designation does nothing to change the role state and local governments have in administering and running elections.

Consider these questions:

- Are you surprised that election infrastructure was not *already* covered?
- Why do you suppose it was not covered before, and for that matter, why do you suppose some "state and local election officials" were "opposed to this designation"?
- The statement lists various things that will *not* happen as a result of the designation. Does the scope of that list give you pause regarding the utility of making this designation in the first place?

Let's use a mini-case study to give you a richer feel for the breadth of the CI definition. Recall that North Korea several years ago engaged in a massive hack of Sony Pictures Entertainment to punish Sony for producing the comedy "The Interview." Was that an attack on America's critical infrastructure, at least insofar as DHS CISA is concerned? To answer questions like that, you have to be familiar with the "sector-specific plans" produced for each of the 16 CI sectors listed above (as required by the "National Infrastructure Protection Plan"). If you read through the sector-specific plan for the "Commercial Facilities" sector, you will find that it includes a variety of more specific "sub-sectors," including the following:

> *Entertainment and Media Subsector . . .*

> The subsector includes media production facilities (e.g., television and motion pictures), print media companies (e.g., newspapers, magazines, and books), and broadcast companies (e.g., television and radio stations). These outlets reach the general population on a continuous basis and have a significant effect on the economy.

Consider these questions:

- Does Sony Pictures Entertainment qualify?
- Should it? Make the case for or against, using the statutory definition of CI quoted earlier.
- Does answering this question one way or the other necessarily dictate whether and how the U.S. government will respond to an attack?

B. The Consequences of a CI Designation

What follows from a CI designation? Not as much as you might expect. One can imagine a world in which this designation carries with it robust authorities for one or more government agencies (CISA being the most likely candidate) to issue regulations compelling or forbidding certain things relating to both cybersecurity and the physical aspects of security. As we have seen, after all, CISA has exactly that sort of authority with respect to cybersecurity for U.S. government entities other than the military and the Intelligence Community. And while it may seem unrealistic as a political matter to grant any regulator such broad authority over the entirety of the private sector, it is not hard to imagine treating all CI-designated entities differently on this dimension.

For the most part, however, that is not how it works. A CI designation as such carries with it no such general regulatory authority. Certain sectors may happen to be subject to regulation (including cybersecurity-related regulation) already thanks to other statutes, but no further authority follows from a CI designation.

Consider these questions:

- Can you explain how the financial sector is an example of the point made in that last sentence? Think back to the GLB Act and the Safeguards Rule.
- What are the best arguments for and against giving CISA broad authority to regulate all CI-designated entities for cybersecurity (and physical security) purposes?
- Is there a good argument for giving such authority to each "sector-specific agency," even if not vesting all of it in CISA alone?
- Should CISA be given such authority as to all entities, period?

Since there is no general authority to regulate all CI-designated entities, the question arises as to what practical consequences follow from the designation. To answer, let's take a close look at excerpts from this Congressional Research Service report on the consequences of the decision to extend CI status to election infrastructure.

The report highlights three "notable consequences" from the designation. First:

> It raised the priority for DHS to provide security assistance to election jurisdictions that request it and for other executive branch actions, such as economic sanctions that the

> Department of the Treasury can impose against foreign actors who attack elements of U.S. CI, including tampering with elections.

Consider these questions:
- Does this mean that such assistance would not be provided otherwise?
- Does this mean that sanctions would not be imposed otherwise?
- Is there real value on this dimension?

Second, the report explained that the designation

> brings the subsector under a 2015 United Nations nonbinding consensus report (A/70/174) stating that nations should not conduct or support cyber-activity that intentionally damages or impairs the operation of CI in providing services to the public. It also states that nations should take steps to protect their own CI from cyberattacks and to assist other nations in protecting their CI and responding to cyberattacks on it. The report was the work of a group of governmental experts from 20 nations, including Russia and the United States.

Consider these questions:
- Do you recognize this reference to the "Group of Government Experts" process that we studied previously?
- Does this point have any practical significance? Be prepared to argue it both ways.

Third, the report explained that the designation also

> provided DHS the authority to establish formal coordination mechanisms for CI sectors and subsectors and to use existing entities to support the security of the subsector. Those mechanisms are used to enhance information sharing within the subsector and to facilitate collaboration within and across subsectors and sectors.

Can you relate this to our prior study of information-sharing mechanisms, including ISACs?

C. The NIST Cybersecurity Framework

Staying with the theme of information-sharing and voluntary improvements to security for privately-owned CI entities, let's turn now to a discussion of the "NIST Cybersecurity Framework."

To understand its role, we need to go back to the 2013. On the same day that President Obama issued PPD-21, he also issued Executive Order 13636. EO 13636 mostly focused on encouraging voluntary information-sharing by CI owner/operators (it was, in that sense, a precursor to CISA 2015). But EO 13636 also sought to encourage better defense of CI entities via this measure:

> *Sec. 7. Baseline Framework to Reduce Cyber Risk to Critical Infrastructure.*
>
> (a) The Secretary of Commerce shall direct the Director of the National Institute of Standards and Technology (the "Director") to lead the development of a framework to

> reduce cyber risks to critical infrastructure (the "Cybersecurity Framework"). The Cybersecurity Framework shall include a set of standards, methodologies, procedures, and processes that align policy, business, and technological approaches to address cyber risks. The Cybersecurity Framework shall incorporate voluntary consensus standards and industry best practices to the fullest extent possible. . . .
>
> (b) The Cybersecurity Framework shall provide a prioritized, flexible, repeatable, performance-based, and cost-effective approach, including information security measures and controls, to help owners and operators of critical infrastructure identify, assess, and manage cyber risk. The Cybersecurity Framework shall focus on identifying cross-sector security standards and guidelines applicable to critical infrastructure. The Cybersecurity Framework will also identify areas for improvement that should be addressed through future collaboration with particular sectors and standards-developing organizations. To enable technical innovation and account for organizational differences, the Cybersecurity Framework will provide guidance that is technology neutral and that enables critical infrastructure sectors to benefit from a competitive market for products and services that meet the standards, methodologies, procedures, and processes developed to address cyber risks. The Cybersecurity Framework shall include guidance for measuring the performance of an entity in implementing the Cybersecurity Framework.

Pursuant to that directive (and consistent with a similar directive from Congress that followed in the Cybersecurity Enhancement Act of 2014), NIST published the first version of the Cybersecurity Framework in February 2014, and then published an updated version in April 2018. The following language (from the original 2014 version) explains what the Framework does—and does not—aspire to do:

> The Framework uses risk management processes to enable organizations to inform and prioritize decisions regarding cybersecurity. It supports recurring risk assessments and validation of business drivers to help organizations select target states for cybersecurity activities that reflect desired outcomes. Thus, the Framework gives organizations the ability to dynamically select and direct improvement in cybersecurity risk management. . . . The Framework provides a common language for understanding, managing, and expressing cybersecurity risk both internally and externally. It can be used to help identify and prioritize actions for reducing cybersecurity risk, and it is a tool for aligning policy, business, and technological approaches to managing that risk. . . . The Core is not a checklist of actions to perform."

Consider these questions:

- How is this similar to MITRE ATT&CK Matrix, and how is it different?
- Does EO 13636 purport to make private sector critical infrastructure owners legally obligated to adopt and adhere to the Cybersecurity Framework?
- Should it do so?
- Is it feasible to create such an obligation via Executive Order, as opposed to legislation?
- Does the existence of the Framework nonetheless cast a legal shadow of sorts, one that might at least create incentives for CI owners? As you ponder that question, think about how CISA 2015 anticipated and addressed parallel concerns.

Another significant part of 2013's EO 13636 is Section 8(e), which states:

> [T]he Secretary of Defense and the Administrator of General Services . . . shall make recommendations to the President . . . on the feasibility, security benefits, and relative merits of incorporating security standards into acquisition planning and contract administration.

Can you relate this to anything we previously studied?

D. Section 9 Entities ("Critical Infrastructure at Greatest Risk")

Up to this point, we have treated all CI equally, even as we defined the category broadly. But of course, CI entities are not all equally significant to the nation; nuclear power plants have more strategic significance than shopping malls, for example. EO 13636 recognized this, carving out a subcategory of particularly important CI that we will refer to as "Section 9 entities" (so named for Section 9 of EO 13636, which establishes the category and which we will examine below). Our goals are to understand the test for qualifying as a Section 9 entity, what consequences follow from that designation, and what the arguments for and against taking a more proscriptive approach might be.

First, let's look at the test for qualifying as a Section 9 entity:

> Sec. 9. Identification of Critical Infrastructure at Greatest Risk.
>
> (a) . . . the Secretary shall use a risk-based approach to identify critical infrastructure where a cybersecurity incident could reasonably result in catastrophic regional or national effects on public health or safety, economic security, or national security.

Applying that test, should the following entities qualify (and if you need more information, what information might that be)?

- nuclear power plant
- hospital that is one of three in a town with 75,000 residents
- shopping mall
- single voting machine

Next, let's have a look at the consequences that follow from receiving this designation. For that, we need to look at EO 13636 Section 10:

> **(a)** Agencies with responsibility for regulating the security of critical infrastructure shall engage in a consultative process with DHS, OMB, and the National Security Staff to review the preliminary Cybersecurity Framework and determine if current cybersecurity regulatory requirements are sufficient given current and projected risks. In making such determination, these agencies shall consider the identification of critical infrastructure required under section 9 of this order. . . .
>
> **(b)** If current regulatory requirements are deemed to be insufficient, . . . agencies identified in subsection (a) of this section shall propose prioritized, risk-based, efficient, and coordinated actions . . . to mitigate cyber risk.

> **(c)** . . . agencies identified in subsection (a) of this section shall, in consultation with owners and operators of critical infrastructure, report to OMB on any critical infrastructure subject to ineffective, conflicting, or excessively burdensome cybersecurity requirements. This report shall describe efforts made by agencies, and make recommendations for further actions, to minimize or eliminate such requirements.
>
> **(d)** The Secretary shall coordinate the provision of technical assistance to agencies identified in subsection (a) of this section on the development of their cybersecurity workforce and programs.
>
> **(e)** Independent regulatory agencies with responsibility for regulating the security of critical infrastructure are encouraged to engage in a consultative process with the Secretary, relevant Sector-Specific Agencies, and other affected parties to consider prioritized actions to mitigate cyber risks for critical infrastructure consistent with their authorities.

Consider these questions:

- Can you sum up the function of each of these subsections?
- Can you think of further steps that *could* have been taken vis-à-vis Section 9 entities, without needing legislation?

It may be tempting to dismiss the combination of Sections 9 and 10 as having little bite. And perhaps that characterization is accurate. But note that it was thought significant enough, at the time, to warrant certain limitations. First, Section 9 itself expressly prohibits applying the "at greatest risk" label to "commercial information technology products" and "consumer information technology services."

What sort of businesses might fall within those categories? Why exempt them?

Separately, Section 9(c) provides that (1) any entity designated as a Section 9 entity shall receive notice of the designation as well as of "the basis for the determination," and (2) DHS must establish a process for which the owner/operators of such entities may then seek reconsideration.

Why?

That was 2013. What has happened since then?

Most recently, in May 2017, President Trump issued Executive Order 13800, Section 2 of which addresses cybersecurity in the CI context. We'll look first at Section 2(b), which is focused on Section 9 entities. In relevant part it provides:

> The Secretary of Homeland Security, in coordination with the Secretary of Defense, the Attorney General, the Director of National Intelligence, the Director of the Federal Bureau of Investigation, the heads of appropriate sector-specific agencies, . . . and all other appropriate agency heads, as identified by the Secretary of Homeland Security, shall:

> **(i)** identify authorities and capabilities that agencies could employ to support the cybersecurity efforts of . . . section 9 entities . . .;
>
> **(ii)** engage section 9 entities and solicit input as appropriate to evaluate whether and how the authorities and capabilities identified pursuant to subsection (b)(i) of this section might be employed to support cybersecurity risk management efforts and any obstacles to doing so;
>
> **(iii)** provide a report to the President . . . that includes . . . findings and recommendations for better supporting the cybersecurity risk management efforts of section 9 entities [based on the above inquiries].

How would you characterize this intervention?

Next, we'll look at Section 2(c), which is not limited to Section 9 entities, but *is* limited to CI entities that are owned by a publicly traded company:

> The Secretary of Homeland Security . . . shall provide a report to the President . . . that examines the sufficiency of existing Federal policies and practices to promote appropriate market transparency of cybersecurity risk management practices by critical infrastructure entities, with a focus on publicly traded critical infrastructure entities.

Can you explain how such an approach might drive improvements to security?

19. Federal Coordination for Significant Cyber Incidents

Notwithstanding all the efforts to boost defense that we have considered, the attacker will get through from time to time. And though most breaches will not be significant enough to warrant a coordinated federal government response, some will. Our final task in this reading is to understand how the federal government has organized itself to identify such situations and what procedures should then be followed.

A. "Significant Cyber Incidents"

The key document here is PPD-41 ("United States Cyber Incident Coordination"), issued by the Obama administration in 2016. Let's start by learning the key categories used in this context, as set forth in PPD-41 Section II ("Definitions"):

> A. **Cyber incident**. An event occurring on or conducted through a computer network that actually or imminently jeopardizes the integrity, confidentiality, or availability of computers, information or communications systems or networks, physical or virtual infrastructure controlled by computers or information systems, or information resident thereon. For purposes of this directive, a cyber incident may include a vulnerability in an information system, system security procedures, internal controls, or implementation that could be exploited by a threat source.

> B. **Significant cyber incident.** A cyber incident that is (or group of related cyber incidents that together are) likely to result in demonstrable harm to the national security interests, foreign relations, or economy of the United States or to the public confidence, civil liberties, or public health and safety of the American people.

Consider these questions:

- What exactly distinguishes a "significant cyber incident" from an ordinary incident?
- Is this definition sufficiently clear so as to yield predictable answers as to which incidents fall into that category?
- Are there good reasons not to have a more precise definition?

B. Lead Agencies for the Three Major "Lines of Effort"

Next, let's turn our attention to PPD-41 Section IV ("Concurrent Lines of Effort"). Again, we are not yet to the part of PPD-41 that addresses interagency coordination, but now we are getting close. Section IV describes three distinct categories of government action that might come into play in the event of any cyber incident (significant or not):

> In responding to any cyber incident, Federal agencies shall undertake three concurrent lines of effort: threat response; asset response; and intelligence support and related activities. . . .
>
> A. Threat response activities include conducting appropriate law enforcement and national security investigative activity at the affected entity's site; collecting evidence and gathering intelligence; providing attribution; linking related incidents; identifying additional affected entities; identifying threat pursuit and disruption opportunities; developing and executing courses of action to mitigate the immediate threat; and facilitating information sharing and operational coordination with asset response.
>
> B. Asset response activities include furnishing technical assistance to affected entities to protect their assets, mitigate vulnerabilities, and reduce impacts of cyber incidents; identifying other entities that may be at risk and assessing their risk to the same or similar vulnerabilities; assessing potential risks to the sector or region, including potential cascading effects, and developing courses of action to mitigate these risks; facilitating information sharing and operational coordination with threat response; and providing guidance on how best to utilize Federal resources and capabilities in a timely, effective manner to speed recovery. . . .
>
> C. Intelligence support and related activities facilitate the building of situational threat awareness and sharing of related intelligence; the integrated analysis of threat trends and events; the identification of knowledge gaps; and the ability to degrade or mitigate adversary threat capabilities.

PPD-41also designated which federal entities shall be "lead agency" for each of these three lines of effort. Section V explains:

> [T]he following agencies shall serve as Federal lead agencies for the specified line of effort:

> 1. In view of the fact that significant cyber incidents will often involve at least the possibility of a nation-state actor or have some other national security nexus, the Department of Justice, acting through the Federal Bureau of Investigation and the National Cyber Investigative Joint Task Force, shall be the Federal lead agency for threat response activities.
>
> 2. The Department of Homeland Security, acting through the National Cybersecurity and Communications Integration Center, shall be the Federal lead agency for asset response activities.
>
> 3. The Office of the Director of National Intelligence, through the Cyber Threat Intelligence Integration Center, shall be the Federal lead agency for intelligence support and related activities.

Each lead agency is responsible not just for conducting operations of its own in relation to that line of effort, but also for coordination with other affected entities, other government agencies, and with the interagency coordinating bodies that PPD-41 establishes (and which we will examine below).

Consider these questions:

- Is each line of effort limited to actions that are responsive to an ongoing or completed incident, or are some proactive? If the latter, which ones?
- Notice the very last clause quoted above ("and the ability to degrade"). Can you think of some examples that would fall under this heading? And does this complicate the issue of lead-agency responsibility for that category?

C. The Cyber Response Group (Policy Coordination)

Now we can focus on the interagency coordination/deconfliction system that PPD-41 establishes in Section V ("Architecture of Federal Government Response Coordination for Significant Cyber Incidents"). First and foremost, it is important to note that the opening line of Section V refers expressly and only to "significant cyber incidents," thus excluding run-of-the-mill cyber incidents from the coordination mechanisms set forth below.

Is that wise? Be able to argue for and against this limitation.

Next, let's look at the distinct coordinating bodies that PPD-41 established or recognized. First, there is the Cyber Response Group ("CRG"). As set forth in part II.A. of the Annex to PPD-41, CRG's duties are to:

> **i.** Coordinate the development and implementation of the Federal Government's policies, strategies, and procedures for responding to significant cyber incidents;
>
> **ii.** Receive regular updates from the Federal cybersecurity centers and agencies on significant cyber incidents and measures being taken to resolve or respond to those incidents;
>
> **iii.** Resolve issues elevated to it by subordinate bodies [more on this below] . . .;

> **iv.** Collaborate with the Counterterrorism Security Group and Domestic Resilience Group when a cross-disciplinary response to a significant cyber incident is required;
>
> **v.** Identify and consider options for responding to significant cyber incidents, and make recommendations to the Deputies Committee, where higher-level guidance is required . . .; and
>
> **vi.** Consider the policy implications for public messaging in response to significant cyber incidents, and coordinate a communications strategy, as necessary, regarding a significant cyber incident.

Who has a seat this table? Any agency can be invited to participate, but as a general matter the CRG should include:

> senior representatives from the Departments of State, the Treasury, Defense (DOD), Justice (DOJ), Commerce, Energy, Homeland Security (DHS) and its National Protection and Programs Directorate, and the United States Secret Service, the Joint Chiefs of Staff, Office of the Director of National Intelligence, the Federal Bureau of Investigation, the National Cyber Investigative Joint Task Force, the Central Intelligence Agency, and the National Security Agency.

CRG meetings are chaired by the "Special Assistant to the President and Cybersecurity Coordinator . . . or an equivalent successor," and they "shall convene on a regular basis and as needed."

D. Cyber Unified Coordination Groups (Operational Coordination)

But wait, there's more. The CRG is designed to focus on strategy and policy, not necessarily operational details in particular instances. It could, of course, be given that role as well. But that's not the path taken in PPD-41. Instead, PPD-41 assigns the job of coordination at the operational level during specific events to a separate interagency body called a "Cyber Unified Coordination Group" ("UCG").

Unlike the CRG, a UCG is not a standing body. Rather, it is called into being only when specific occasions arise. And what is its function, exactly? PPD-41 Part V explains that "A Cyber Unified Coordination Group (UCG) shall serve as the primary method for coordinating between and among Federal agencies in response to a significant cyber incident as well as for integrating private sector partners into incident response efforts, as appropriate." The Annex elaborates that a UCG shall act to ensure inclusion of all relevant agencies in the response effort, coordinate the planning and execution of both response and recovery tasks, coordinate international and cross-sector outreach, facilitate information-sharing, and coordinate communications with impacted parties as well as the public. If the incident has physical effects, moreover, the UCG is to combine its efforts with the relevant lead agency or interagency group managing those consequences.

Who has a seat at the table with a Cyber UCG? PPD-41 specifies that the group should include each of the lead agencies for threat response, asset response, and intelligence support, along with any relevant SSA. At discretion, the UCG also may include other federal agencies, agencies at the state/local/tribal levels, NGOs, international partners, and private-sector entities.

Who decides when to invite such additional participants, and for that matter, who decides when a Cyber UCG is warranted in the first place? PPD-41 explains:

> A Cyber UCG shall be formed at the direction of the NSC Principals Committee, Deputies Committee, or the CRG, or when two or more Federal agencies that generally participate in the CRG, including relevant SSAs, request its formation. A Cyber UCG shall also be formed when a significant cyber incident affects critical infrastructure owners and operators identified by the Secretary of Homeland Security as owning or operating critical infrastructure for which a cyber incident could reasonably result in catastrophic regional or national effects on public health or safety, economic security, or national security.

Can you think of a scenario in which the "catastrophic effect" test is met, yet none of the listed entities have asked for formation of a Cyber UCG? If your answer is no, what is the point of having that additional pathway?

PPD-41 includes an important caveat:

> The Cyber UCG is intended to result in unity of effort and not to alter agency authorities or leadership, oversight, or command responsibilities. Unless mutually agreed upon between agency heads or their designees, and consistent with applicable legal authorities such as the Economy Act of 1932 (31 U.S.C. 1535), Federal departments and agencies will maintain operational control over their respective agency assets.

Can you anticipate problems that might arise under this heading?

E. The National Cyber Incident Response Plan

The last part of the Annex to PPD-41 directed DHS to work with others to create a "national cyber incident response plan" within six months. In December 2016, DHS accordingly published The National Cyber Incident Response Plan ("NCIRP"). Think of the NCIRP as the detailed playbook that implements the vision spelled out in PPD-41.

The full version is available here if this interests you, but I only need you to look at one element, which appears in full below.

The NCIRP tries to help us resolve a critical threshold question: Are we dealing with a "significant cyber incident" in the first place, thus triggering the CRG and UCG systems?

The answer the NCIRP gives is found on p. 29: "Cyber incidents rated a '3' or greater" on the NCIRP's schema for this question "will equate to a significant cyber incident." Here's the schema (found in Annex B to the NCIRP):

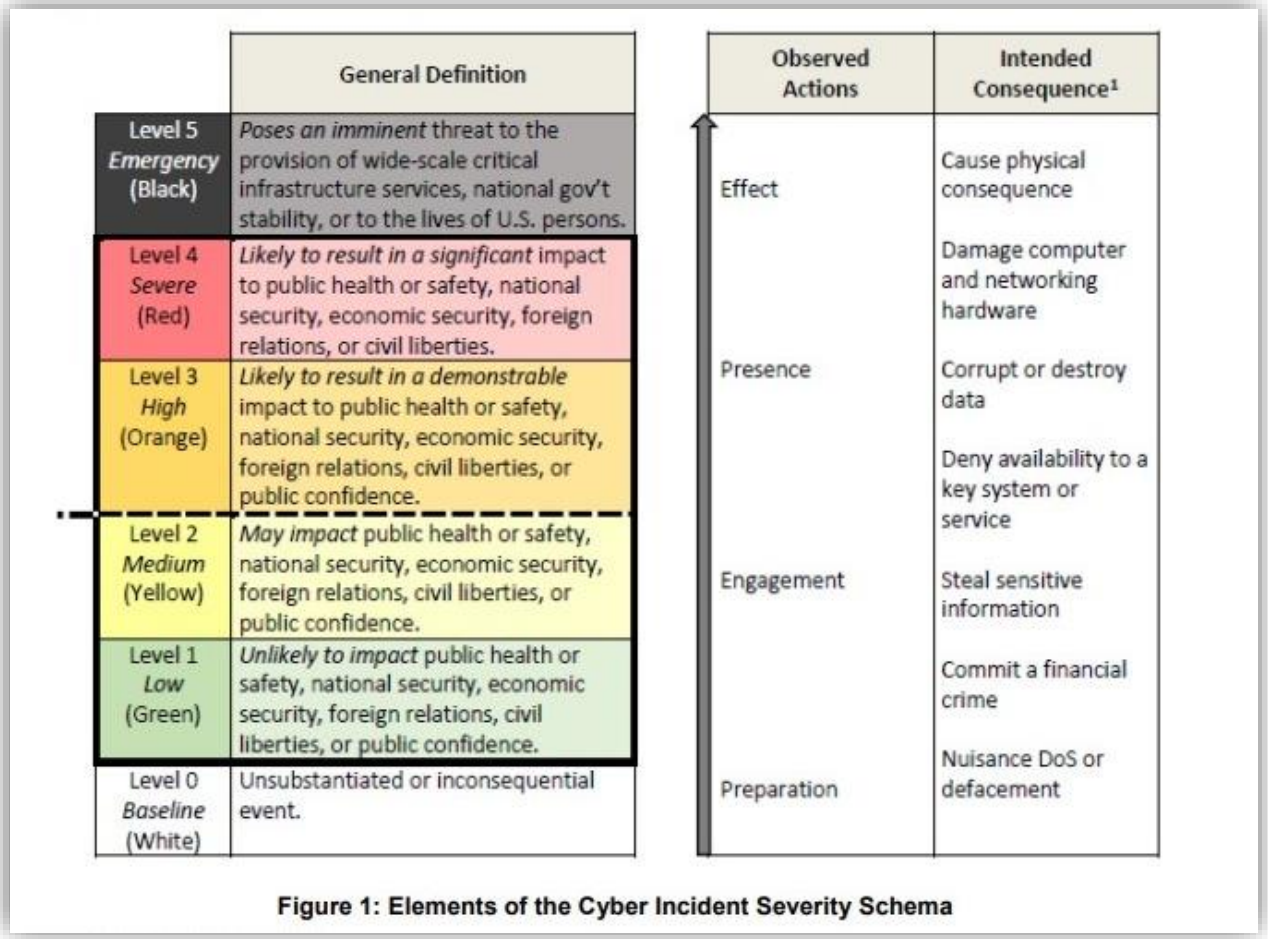

	General Definition
Level 5 *Emergency* (Black)	*Poses an imminent* threat to the provision of wide-scale critical infrastructure services, national gov't stability, or to the lives of U.S. persons.
Level 4 *Severe* (Red)	*Likely to result in a significant* impact to public health or safety, national security, economic security, foreign relations, or civil liberties.
Level 3 *High* (Orange)	*Likely to result in a demonstrable* impact to public health or safety, national security, economic security, foreign relations, civil liberties, or public confidence.
Level 2 *Medium* (Yellow)	*May impact* public health or safety, national security, economic security, foreign relations, civil liberties, or public confidence.
Level 1 *Low* (Green)	*Unlikely to impact* public health or safety, national security, economic security, foreign relations, civil liberties, or public confidence.
Level 0 *Baseline* (White)	Unsubstantiated or inconsequential event.

Observed Actions	Intended Consequence[1]
Effect	Cause physical consequence
	Damage computer and networking hardware
Presence	Corrupt or destroy data
	Deny availability to a key system or service
Engagement	Steal sensitive information
	Commit a financial crime
Preparation	Nuisance DoS or defacement

Figure 1: Elements of the Cyber Incident Severity Schema

Consider these questions:

- Do the "general definitions" in this chart seem likely to truly help you make a categorization determination?
- Does the "observed action" column help?
- Does the "intended consequence" column help?
- Can you explain how "intended consequence" differs from "observed" or "actual" consequences, and is there any problem with using "intended" as a test?

II. THE OFFENSIVE PERSPECTIVE

We've been defense-focused up to this point, in the specific sense of surveying the institutions, laws, and policies intended to deter, punish, and mitigate the consequences of unauthorized access. For better or worse, however, limiting unauthorized access is not always the ultimate policy goal. In the U.S. system, some institutions, laws, and policies promote (or at least tolerate) *offense*—that is, efforts to penetrate or interfere with a system without its owner's authorization (or, perhaps, awareness). For the sake of convenience, we might call this lawful-but-unauthorized access (meaning lawful from a *U.S.* perspective; needless to say, such activity may well violate the laws of other countries when they occur overseas).

You will note immediately, I hope, that the very idea that this category exists appears at first blush to be in considerable tension with the policy goals advanced by, well, pretty much everything we studied in Unit I. Why, then, should there even be such a category? We will explore that question across several contexts.

To a substantial extent, lawful-but-unauthorized access frameworks rest on the counterintuitive claim that they can, in the right circumstances, promote security. We see this, for example, in the "hackback" scenario, in which some advocate empowering the private sector to respond to an attack with self-help measures that will have effects outside their own networks (that is, effects on the attacker's network or, more likely, effects on intermediary networks from which an attack has been staged). And we see this as well with the portion of U.S. Cyber Command's mission that entails out-of-network operations in "gray space" and "red space" conducted with the intent to identify, disrupt, or prevent adversary efforts to hack U.S. or allied systems. It is in this sense that Unit I opened with a note regarding the affirmative use of unauthorized access methods to promote "defense."

But the case for lawful-but-unauthorized access does not have to rest entirely on that defensive-minded ground. In some cases, in fact, lawful-but-unauthorized access frameworks are intended to serve goals that are simply distinct from cybersecurity. We see this with law enforcement investigations intended to solve (or prevent) crimes of all sorts; with collection of foreign intelligence (that is, espionage) on a vast array of topics; with promotion of U.S. foreign policy goals via covert action; and with military operations both within and outside the context of armed conflict.

Our aim in this unit is to survey each of those scenarios, with an emphasis on the key institutions, policy conflicts, and legal frameworks.

20. Lawful Private Sector Hacking?

Are there circumstances in which we want someone in the private sector to be able to access another's system without their permission? We just completed a long unit focused on how the United States discourages that sort of thing—including by imposing criminal and civil liability under the CFAA—so the idea of instead tolerating or even encouraging it at first blush seems jarring. The issue proves to be complicated, however, when we ask this question in the specific context of a hacking victim hoping to respond in self-defense (*i.e.*, "hackback").

In recent years there has been considerable debate regarding both the extent to which certain defensive measures are permissible in light of the CFAA (as well as CISA 2015's "defensive measures" provision), and whether and how it might be desirable to prune back the CFAA in order to remove disincentives for potential or actual hacking victims to protect themselves in new (or at least newly legal) ways.

A. Defining Terms

Before we proceed, we should note that there is some disagreement regarding the proper meaning of the term "hackback." Some use it to refer to any defensive measure adopted by a potential victim where the measure will have a downstream effect inside an adversary's system. On that view, even a simple beacon would count. Others would reserve the term for more aggressive forms of self-defense, limiting it to measures that have a disruptive effect on the adversary's system or that provide the victim (or whomever is assisting the victim, for the victim may turn to an outside security vendor for help) with ongoing unauthorized access to some part of that system. Those who take this narrower view sometimes distinguish between "hackback" and "active defense," with the latter referring to less-aggressive measures (like beacons) that are not disruptive and that do not provide ongoing access. On the other hand, you sometimes will see the phrase "active defense" used more broadly to span across this entire category. The important point, at any rate, is that you should make sure to make clear how *you* are using these labels, and likewise that you understand how others use them.

For the sake of convenience, I will use "hackback" as a catch-all phrase meant to encompass the entire category of self-defense measures that might have an effect within the attacker's own system.

Now we can proceed. Let's first identify the competing policy considerations at work in the hackback debate, and then move on to consider the challenges that arise once one decides to prune the CFAA in hopes of encouraging at least some forms of hackback.

B. Competing Policy Considerations

Let's use a hypothetical scenario in order to bring to light the competing policy considerations that drive the debate over hackback.

Assume an OPM-like scenario involving a private-sector business which we will call Cuckoo's Eggnog Company (or simply "Cuckoo"). You are Cliff, the CEO.

Cuckoo's Chief Information Security Officer ("CISO") has just notified you of an unfolding incident:

> "Bad news, Cliff. Somebody's inside our network. Looks like they phished the creds of someone on my team, and they've been logging in remotely at night with admin privileges. We're still sorting out what they've been up to, but so far know that they've been scouting our files, and clearly took an interest in our secret formulas for next year's new products. Anyway, looks like our uninvited guest copied some key files, including that database we created regarding customer preferences. He or she got it all stored in a particular spot under an innocuous name, and compressed with an eye towards a quiet, gradual exfiltration. And, well, looks like they went ahead and extracted it this morning. But it's not all bad news. We're pretty sure we have the IP address of the server they sent those files to. And…well, you know…we've got some smart folks on our team. I think we could try some…creative things to get our stuff back. If you want to."

By this time, your general counsel has joined the meeting. He has a nervous look on his face. "What exactly do you have in mind?" he asks the CISO.

"Well, for starters I'd like to do some scanning, see if maybe this server they're using has some known vulns we might be able to use."

"Use for what, exactly?"

"That's up to y'all. But, strictly hypothetically speaking, maybe I could delete their copy of our data and files."

You perk up. "You can do that?"

Encouraged by your interest, the CISO leans in. "Oh yeah, for sure. Or I could leave it all there, but stick in an extra file, something with a tempting title but a little surprise. They'll eventually open it, and when they do it'll phone home, giving me an IP address that would help us know if they move these files somewhere else."

You are definitely tempted. Before you can answer, though, she continues: "But you know what would be even better, though? Let's pop them with some ransomware. Or maybe just delete everything on the whole system!"

Your CISO is smiling now, but your general counsel is not.

"Whoa, slow down! I think we should call the cops, and leave this to the pros."

The CISO rolls her eyes. "Oh right, yes, that will be super helpful. Give me a break. You don't know who to call, and even if you did it would be too late. Let me fix this while we still can. I can do it right now."

Consider the following:

- Make a list of the distinct actions the CISO has suggested.
- For each one, what exact benefits might it provide? Think in terms of both the company and society.
- Is it likely that government action could produce that same benefit?
- If you *were* to seek government help, whom would you call?
- Now identify potential negative consequences, apart from legal liability.
- Does the balance of pros and cons (again setting aside legality) favor taking any of these steps?
- Now assess the legality of each of these steps under the CFAA (bearing in mind, too, the "defensive measures" provision in CISA 2015).

Note that the manual on computer crimes published by DOJ CCIPS includes an express admonition discouraging victims of a hack from responding in kind:

> Do Not Hack into or Damage the Source Computer
>
> Although it may be tempting to do so (especially if the attack is ongoing), the company should not take any offensive measures on its own, such as "hacking back" into the attacker's computer—even if such measures could in theory be characterized as "defensive." Doing so may be illegal, regardless of the motive. Further, as most attacks are launched from compromised systems of unwitting third parties, "hacking back" can damage the system of another innocent party. If appropriate, however, the company's system administrator can contact the system administrator from the attacking computer to request assistance in stopping the attack or in determining its true point of origin.

Does this change your analysis above?

C. Pruning the CFAA to Pave the Way for Hackback?

Some members of Congress (led by Rep. Tom Graves of Georgia) have proposed legislation to prune back the CFAA, with an eye toward encouraging certain hackback-related measures. A bill known as the "Active Cyber Defense Certainty Act"—that is, the "AC/DC Act"—languished in the last session of Congress, and a new version with the same name (filed in June 2019) currently shows no signs of moving forward. Nonetheless, it is instructive to look closely at the AC/DC Act because it illustrates the difficulties involved in amending federal law to create more space for such self-defense measures.

The full text is available here, but the excerpts you need to read all appear below. The AC/DC Act would prune back the CFAA in two distinct ways. First, it would more clearly establish the legality of using "attributional technologies." Second, it would bless a category of activity it calls "active cyber defense measures." Let's take those in order.

1. "Attributional Technology"

Section 3 of the bill aims to ensure the legality of using beacons, among other "attributional technologies." To achieve that end, it would amend the CFAA to include this language:

> Exception For The Use Of Attributional Technology.—
>
> **(1)** [Section 1030(a)] shall not apply with respect to the use of attributional technology in regard to a defender who uses a program, code, or command for attributional purposes that beacons or returns locational or attributional data in response to a cyber intrusion in order to identify the source of an intrusion; if—
>
> **(A)** the program, code, or command originated on the computer of the defender but is copied or removed by an unauthorized user; and
>
> **(B)** the program, code, or command does not result in the destruction of data or result in an impairment of the essential operating functionality of the attacker's computer system, or intentionally create a backdoor enabling intrusive access into the attacker's computer system.

What are the pros and cons of this approach?

Note that Section 3 also offers a definition of "attributional data":

> The term 'attributional data' means any digital information such as log files, text strings, time stamps, malware samples, identifiers such as user names and Internet Protocol addresses and metadata or other digital artifacts gathered through forensic analysis."

Is that too narrow, too broad, or just right?

Of course, hacking back can involve more than just beacons and the like. What about more aggressive measures?

2. "Active Cyber Defense Measures"

Section 4 of the bill addresses a category of activity labeled "active cyber defense measures" (ACDMs). ACDMs are defined to include, as a default matter, any measure that is:

> **(I)** undertaken by, or at the direction of, a defender; and
>
> **(II)** consist[s] of accessing without authorization the computer of the attacker to the defender's own network to gather information in order to—
>
> **(aa)** establish attribution of criminal activity to share with law enforcement and other United States Government agencies responsible for cybersecurity;
>
> **(bb)** disrupt continued unauthorized activity against the defender's own network; or
>
> **(cc)** monitor the behavior of an attacker to assist in developing future intrusion prevention or cyber defense techniques; but

Consider these questions:

- In what ways is this broader than Section 3's "attributional technology" concept?
- Is that breadth potentially problematic? (As we will see below, the statute goes on to carve out limits to this definition; those limits will make more sense, though, if you ponder the issues that would arise without them.)

As noted, Section 4's broad definition of the ACDM category is subject to a number of specific carveouts. An activity cannot count as an ACDM after if it falls into any of the following seven categories:

> **(I)** intentionally destroys or renders inoperable information that does not belong to the victim that is stored on another person or entity's computer;

Consider these questions:

- Does this foreclose deletion of files copied from the victim's system?
- What about files copied from the system of some other victim?
- What if the victim/defender in good faith believes a file is a copy of its own data and hence deletes it, but this turns to be mistaken?

> **(II)** recklessly causes physical injury or financial loss . . .;

Is this a proper limitation? Why not limit the statute to intentional harms? Or expand it to include negligent harms?

> **(III)** creates a threat to the public health or safety;

Can you define the following terms?

- Threat
- Public health
- Public safety

> **(IV)** intentionally exceeds the level of activity required to perform reconnaissance on an intermediary computer to allow for attribution of the origin of the persistent cyber intrusion;

How will the victim/defender know that the system in question is an "intermediary"?

> **(V)** intentionally results in intrusive or remote access into an intermediary's computer;

Is this exception consistent with the bill's effort to permit beacons or even more aggressive forms of active defense?

> **(VI)** intentionally results in the persistent disruption to a person or entity's internet connectivity resulting in damages . . .; or

Consider these questions:

- Why have an exception for this scenario?
- How do we know when disruption becomes "persistent"?

> **(VII)** impacts any computer described under subsection (a)(1) regarding access to national security information, subsection (a)(3) regarding government computers, or to subsection (c)(4)(A)(i)(V) regarding a computer system used by or for a Government entity for the furtherance of the administration of justice, national defense, or national security;

Why have this exception?

Now that we have a handle on what sorts of activities might count under Section 4 as ACDMs, the question becomes: So what? The answer is that Section 4 would eliminate *criminal* liability—and *only* criminal liability—under the CFAA for qualifying ACDMs:

> **(1)** GENERALLY.—It is a defense to a criminal prosecution under this section that the conduct constituting the offense was an active cyber defense measure.
>
> **(2)** INAPPLICABILITY TO CIVIL ACTION.—The defense against prosecution created by this section does not prevent a United States person or entity who is targeted by an active defense measure from seeking a civil remedy, including compensatory damages or injunctive relief pursuant to subsection (g).

Why remove criminal, but not civil liability risk?

Notice that the precise effect of Section 4 is to give an affirmative defense to a qualifying victim/defender, not to directly remove CFAA liability as such.

Does this matter?

Another wrinkle: To be a "defender" eligible for the ACDM defense, Section 4 specifies that the victim must be "a victim of a persistent unauthorized intrusion of the individual entity's computer."

Consider these questions:

- Why limit the victim category to those who suffer "persistent" intrusion?
- How many instances must there be, and of what duration, in order to satisfy this test?
- Why exclude one-off intrusions?

There's more. In order to get the benefit of the ACDM defense, a qualifying "defender" also must follow a special notification regime. Section 5 specifies that a

> defender who uses an active cyber defense measure under the preceding section must notify the FBI National Cyber Investigative Joint Task Force and receive a response from the FBI acknowledging receipt of the notification prior to using the measure. . . . Notification must include the type of cyber breach that the person or entity was a victim of, the intended target of the active cyber defense measure, the steps the defender plans to take to preserve evidence of the attacker's criminal cyber intrusion, as well as the steps they plan to prevent damage to intermediary computers not under the ownership of the attacker and other information requested by the FBI to assist with oversight.

Consider these questions:

- What is the point of requiring this?
- Does the prospect of delay defeat the purpose of creating a pathway for private use of ACDMs?
- What sort of burdens does this place on FBI's NCIJTF?
- Does this notification regime create any risk of unintended consequences, from the point of view of a foreign government who is aware of the notification system and then detects a U.S.-based, private-sector actor using ACDMs against it?

21. The Insecurity Industry

Let's turn now to a related topic involving "legitimate" private-sector hacking: the policy and legal questions associated with what some describe pejoratively as the "insecurity industry."

A. The Insecurity Industry, Government Clients, and the Vulnerabilities Equities Process Issue

In this course, we have spoken often of white-hat security researchers, corporate ventures like Project Zero, bug-bounty programs, and other elements of the security-enhancing ecosystem. These players differ in whether and how they seek to be paid for their efforts, but they have in common a general commitment to increasing information security for all by using their knowledge of vulnerabilities, exploits, and tactics in a way that's intended to cause the owner or creator of relevant systems to patch or take other remedial action. We also have spoken of actors who use the same knowledge for nefarious purposes, either employing it directly on their own or cashing in on the black market by selling products and services without regard for how they might be put to use.

Between these poles, we find the insecurity industry. The key to understanding the insecurity industry is to grasp that, for better or worse, most domestic legal systems permit certain entities (usually limited to government agencies, such as law enforcement organizations) to engage in unauthorized access under at least some conditions. That being the case, those entities either must develop their own capacities to execute such operations or else purchase that capacity from others. And that's where the insecurity industry steps in: these entities supply products and services to those legitimate users. Or that's how it is supposed to work, at any rate.

Unquestionably, it does sometimes work that way. A famous example occurred several years ago, when the FBI had warrants to search the contents of iPhones belonging to Syed Farook and Tashfeen Malik, the perpetrators of a horrific terrorist attack in San Bernadino in December 2015. The FBI did not have their passwords, however, and could not run a brute-force password-guessing program without running afoul of iPhone security features that would progressively slow down the time between guesses or even wipe the data from the phone after a certain number of incorrect guesses. The FBI turned to Apple for help, and Apple cooperated with the FBI up to a point. But, critically, it refused a request to develop a bespoke iOS update that would have negated the security features. The FBI then took Apple to a court of law, and Apple in turn took the FBI to the court of public opinion. Intense controversy followed, until the FBI announced that an unidentified third-party somehow had managed to solve the access challenge. Here's how Ellen Nakashima described this turn of events for the Washington Post:

> The FBI cracked a San Bernardino terrorist's phone with the help of professional hackers who discovered and brought to the bureau at least one previously unknown software flaw, according to people familiar with the matter. The new information was then used to create a piece of hardware that helped the FBI to crack the iPhone's four-digit personal identification number without triggering a security feature that would have erased all the data, the individuals said. The researchers, who typically keep a low profile, specialize in hunting for vulnerabilities in software and then in some cases selling them to the U.S. government. They were paid a one-time flat fee for the solution.

The public later learned that the FBI paid approximately $900,000 for the capability it acquired in that one instance. The Nakashima article goes on to explain:

> FBI Director James B. Comey has said that the solution works only on iPhone 5Cs running the iOS 9 operating system — what he calls a "narrow slice" of phones.
>
> Apple said last week that it would not sue the government to gain access to the solution. Still, many security and privacy experts have been calling on the government to disclose the vulnerability data to Apple so that the firm can patch it.
>
> If the government shares data on the flaws with Apple, "they're going to fix it and then we're back where we started from," Comey said last week in a discussion at Ohio's Kenyon College. Nonetheless, he said Monday in Miami, "we're considering whether to make that disclosure or not."
>
> The White House has established a process [the "Vulnerabilities Equities Policy and Process," or "VEP"] in which federal officials weigh whether to disclose any security vulnerabilities they find. . . . The policy calls for a flaw to be submitted to the process for consideration if it is "newly discovered and not publicly known." "When we discover these vulnerabilities, there's a very strong bias towards disclosure," White House cybersecurity coordinator Michael Daniel said in an October 2014 interview. . . . "That's for a good reason. If you had to pick the economy and the government that is most dependent on a digital infrastructure, that would be the United States." But, he added, "we do have an intelligence and national security mission that we have to carry out. That is a factor that we weigh in making our decisions."
>
> The decision-makers, which include senior officials from the Justice Department, FBI, National Security Agency, CIA, State Department and Department of Homeland Security, consider how widely used the software in question is. They also look at the utility of the flaw that has been discovered. Can it be used to track members of a terrorist group, to prevent a cyberattack, to identify a nuclear weapons proliferator? Is there another way to obtain the information?
>
> In the case of the phone used by the San Bernardino terrorist, "you could make the justification on both national security and on law enforcement grounds because of the potential use by terrorists and other national security concerns," said a senior administration official, speaking on the condition of anonymity because of the matter's sensitivity.
>
> A decision also can be made to disclose the flaw — just not right away. An agency might say it needs the vulnerability for only a few months. "A decision to withhold a vulnerability is not a forever decision," Daniel said in the earlier interview. "We require periodic reviews. So if the conditions change, if what was originally a true [undiscovered flaw] suddenly becomes identified, we can make the decision to disclose it at that point."

Note: If you wish, you can read more about the "Vulnerabilities Equities Policy & Process" here, though you do not need to so for purposes of this class.

Consider the following:

- Can you explain the relationship between the growing ubiquity of encryption both on devices themselves and on communications-in-transit, and the idea that the government is motivated to devote more resources to acquiring hacking capabilities?
- Be prepared to attack or defend the idea of the "Vulnerabilities Equities Policies and Process," in accordance with your own view of the matter.
- Be prepared to make policy arguments for and against allowing the government to buy hacking capabilities from outsiders in the specific context in which a law enforcement agency has obtained a warrant authorizing the search.
- Does your answer apply the same way as to all countries and their varied legal systems, or is it a U.S.-specific answer?
- Does your answer change if the government does not have a warrant?
    - What if the situation is a foreign-intelligence investigation rather than a criminal investigation?
    - What if the government's intent is to use the capability only against foreign targets outside the United States?
- If such sales are made unlawful, what impact would that have on the related debate over whether the law should empower courts under certain conditions to order hardware and software vendors like Apple to cooperate with the government in developing solutions to circumvent user-privacy features?

B. Sales that Facilitate Abuse

If every insecurity-industry vendor sold only to the FBI in contexts involving warrants, this topic would be much less interesting. In reality, though, there are companies around the globe selling such goods and services to a wide variety of government buyers, with some of their aims laudable and others not so at all. To understand this better, we have an excerpt from an article that Joseph Cox and Lorenzo Franceschi-Bicchierai wrote for Motherboard in 2018:

> At first glance, Azimuth Security looks like any other bustling startup. Photos tweeted by the firm's co-founder show a staffer zipping in front of glass-walled conference rooms on a hoverboard and employees in T-shirts playing with a stylish chess set over a beer. But this small Australian company plays a crucial role in the continuous battle for spies and cops to hack into phones around the world. . . .
>
> The story of this little-known company provides a rare peek inside the secretive exploit trade, which is populated with military contractors, individual researchers, and boutique high-end hacking shops like Azimuth. While the trade is commonly painted as a wild west full of mercenaries who sell hacking tools to whoever can afford them, over a dozen well-placed sources described an overlooked section of the industry that focuses on supplying to a select group of democratic governments, rather than authoritarian regimes. . . . Three sources familiar with the company said Azimuth—through its partner firm—provides exploits to members of the so-called Five Eyes, a global intelligence sharing group made up of the United States, United Kingdom, Canada, Australia, and New Zealand. . . . Azimuth's exploits are used in terrorism cases, and potentially other types of crime such as kidnapping or child pornography as well. . . .

> A plethora of companies now focus solely on offering exploits and other hacking tools to intelligence and law enforcement agencies around the world, typically with customer support and additional products to extract information from target devices. Think of it as a hacking-tools-as-a-service. . . . Many of these firms have controversial client lists, including countries with abysmal human rights records such as Sudan, Ethiopia, and Russia.

Consider the following:

- What distinction does this passage draw in terms of the potential purchasers of these "hacking tools"?
- Should that distinction matter with reference to the legality of such sales?
- If you think purchaser identity should be used in determining the legality of such sales, what statutory language would suffice to implement the distinction?
- The passage also highlights various investigative purposes the purchaser might have in mind. Should that variable be used to determine legality?
- If so, what statutory language would suffice to implement the distinction?
- To what extent is it reasonable to expect a seller to determine, accurately, who will be the ultimate user of the capabilities it sells and the purposes to which the capabilities will be put?

It is clear that, on some occasions, vendors have sold hacking tools to an entity that has, in turn, used them to facilitate some undesirable end, such as surveilling political opponents (or worse). What liability, if any, should attach to the vendor in such cases?

Consider the suit Facebook/WhatsApp recently filed against the Israeli firm NSO Group. Here are excerpts from a handy summary by Andy Greenberg at Wired:

> WhatsApp published a statement accusing NSO of targeting 1,400 of its users, including at least 100 members of "civil society" such as journalists and human-rights defenders, with malicious voice calls designed to infect targeted phones with malware and steal messages despite WhatsApp's end-to-end encryption. Those numbers would represent a new scale for NSO, whose malware has already been linked to attacks against activists ranging from the now-imprisoned United Arab Emirates dissident Ahmed Mansoor to Mexican activists opposing a soda tax.
>
> WhatsApp paired its statement with a lawsuit in a Ninth Circuit court, accusing NSO of violating the Computer Fraud and Abuse Act. . . . To make that charge stick, WhatsApp will have to show that NSO obtained illegal access to WhatsApp's own systems. Given that NSO's targets were WhatsApp users rather than, say, WhatsApp's servers, they'll have to find an argument that they, as the plaintiff, were the victim. . . .
>
> WhatsApp's most obvious unauthorized access argument relates to its terms of service, which prohibit reverse-engineering WhatsApp's code, harming its users, or sending malware via WhatsApp. The company might argue that by agreeing to those terms of service and then violating them, NSO's use of WhatsApp was unauthorized all along. The complaint appears to lay the groundwork for that case: It points out that NSO Group staff "created various WhatsApp accounts and agreed to the WhatsApp Terms."
>
> But that terms-of-service argument will be an uphill battle . . . the Ninth Circuit in particular has set a clear precedent that terms-of-service violations alone don't constitute unauthorized access. . . . WhatsApp's lawsuit doesn't make any mention of prior notice to NSO to stop abusing its services or hacking its users. . . .

> Another, trickier strategy for WhatsApp may be to claim that the malicious data NSO sent via WhatsApp servers *was itself* a kind of unauthorized access. The WhatsApp complaint accuses NSO of initiating malicious calls that hid their attack code in fake settings data, and in doing so bypassed "technical restrictions" on what sort of data WhatsApp's servers were designed to pass on to phones. This may be the crux of WhatsApp's CFAA claim: that WhatsApp's own access restrictions were "hacked" with this technique, just as if someone had bypassed a more obvious access restriction like one that demanded a username and password. . . .
>
> "We dispute today's allegations and will vigorously fight them," said NSO in a statement. "The sole purpose of NSO is to provide technology to licensed government intelligence and law enforcement agencies to help them fight terrorism and serious crime. Our technology is not designed or licensed for use against human rights activists and journalists."

Consider the following questions:

- Can you explain how the first theory of CFAA liability in this case relates to the *Facebook v. Power Ventures* case?
- If Facebook/WhatsApp prevail on this theory, does this have implications for the insecurity industry more generally?
    - For example, would this call into question the legality of the capability the FBI apparently purchased in order to access the San Bernadino iPhones?
- Can you explain the second theory of CFAA liability in this case?
    - If it prevails, would it cast a significant shadow over other insecurity-industry practices?

C. Using Export-Control Laws and Institutions to Limit Risky International Sales

In addition to the question of what is and should be legal in any given country's domestic legal system, there is a distinct set of questions concerning whether and how the domestic and international legal systems that regulate international trade—particularly trade in armaments and other security-sensitive capabilities—apply to the insecurity industry. This has been a hot topic in recent years in light of periodic stories of firms in one country providing capabilities to foreign governments which then put them to abusive ends.

A prominent example involves the Italian firm "Hacking Team," which in 2015 was itself subject to a massive doxing episode that revealed to the public that the company sold its surveillance-enabling malware kits to a variety of regimes including Azerbaijan, Bahrain, Egypt, Ethiopia, Kazakhstan, Morocco, Nigeria, Oman, Saudi Arabia, and Sudan, but also the DEA, the FBI, and DoD in the United States. This added fuel to the fire of an ongoing effort to update international arms-trade rules to encompass such transactions. Andy Greenberg at Wired summed up the situation at the time:

> That issue of whether hacking tools are defined as weapons in the terms of arms control agreements couldn't be more timely: An arms control pact called the Wassenaar Arrangement has been hotly debated in recent weeks over its measures that would control the international export of intrusion software. . . . The Wassenaar Arrangement has been criticized by the hacker community as limiting security research and preventing the sharing of penetration testing tools. But Privacy International's Eric King argues that the practices of Hacking Team demonstrate why the pact is necessary, along with what he

> describes as "carve-outs" to protect security research. "What's clear is that these companies can't be left to their own devices," says King. "Some form of regulation is needed to prevent these companies from selling to human rights abusers. That's a hard policy question, and one tool won't be a silver bullet. But regulation and export controls should be part of the policy response."

Consider the following questions:

- Should export-control rules encompass vulns, malware, or both?
- Why would "White Hat" security researchers object?

What happened next? Garret Hink explained here for Lawfare in 2018:

> The United States successfully negotiated research-use exceptions to export controls on surveillance tools at the December 2017 meeting of the Wassenaar Arrangement, a club of advanced economies that coordinates export controls. These export controls—requirements that organizations selling or sending technologies with potential military applications abroad obtain a license from the Commerce Department—affect key swaths of the cybersecurity industry. Although countries implement export controls at the national level, the United States and 40 other countries have agreed to coordinate their controlled items at the Wassenaar Arrangement, an international framework for creating a voluntary export control regime. At this year's meeting, the U.S. aimed to correct what the cybersecurity industry portrays as overly-broad controls on intrusive surveillance software—controls that security experts say "criminalized" essential tools for stopping malware. After years of debate over the proper scope of export controls on surveillance products, the U.S. has finally made a beachhead on getting long-sought-after exemptions for security research and information sharing. In this post, I describe the original Wassenaar export controls, summarize the 2017 revisions, and forecast what we should expect to see next.
>
> **Background**
> The Arab Spring revealed how repressive regimes use Western commercially developed software surveillance tools to spy on dissidents and human rights activists. Human rights organizations sued a French company for giving to the Libyan government equipment that activists say enabled torture of dissidents. Privacy International and other civil-society groups pressured the British government to use existing legal mechanisms restrict repressive regimes' access to network intrusion software that employed enabled governments to intercept email, instant messaging and webcam data. (Citizen Lab's research explores this topic extensively.) In 2013, the British and French governments negotiated the addition of two types of dual-use technology—"intrusion software" and "IP network communications surveillance systems"—to the lists of dual-use technologies that the Wassenaar Arrangement governs
>
> The Wassenaar Arrangement is an export control framework—not an international regulatory agency or treaty organization, but rather, a group of countries that meet regularly and agree to control certain technologies. An export control is a requirement that a company wishing to sell a product abroad get a government license to export the item; it is not a ban on that item's export. Wassenaar has no way to make its controls legally binding on its members, who regulate controlled items through their domestic export control regimes. The arrangement's 41 members include the U.S., near-all the European Union (Cyprus is the lone outlier), Russia, Turkey, Argentina and South Africa. Its goal is to prevent "destabilising accumulations" of conventional arms and dual-use goods and technologies—items with both civilian and military applications.

The Intrusion Software Controls

The "intrusion software" control took on the difficult task of regulating surveillance software based on computer code functionality. Wassenaar defined intrusion software as "software specially designed or modified to avoid detection by monitoring tools, or to defeat protective countermeasures" and that either extracted data from a computer or network device or modified the "standard execution path" of a program to allow "the execution of externally provided instructions."

But rather than control intrusion software itself, the arrangement put export controls on software, systems or equipment that interacted with intrusion software. The provision would cover the software toolkits that companies sell to law enforcement and intelligence agencies to carry out intrusive surveillance—see for example Hacking Team's notorious RCS package. It also controlled any type of technology involved in the development of intrusion software. "Technology" in the control's context meant essentially any program, code or software tool that was connected to intrusion software. In one interpretation of the WA's controls, "intrusion software" meant code that took advantage of an exploit. By controlling software that used software vulnerabilities to carry out surveillance, the WA control targeted a limited subset of items that related to software exploits.

When the United States tried to implement the intrusion software controls, Symantec, FireEye, independent security researchers and the EFF raised serious concerns about their effect on security software and research. First, the Commerce Department's implementation through the Commerce Department's Bureau of Industry and Security (BIS) expanded its scope to cover a broad range of cybersecurity items. Mailyn Fidler detailed for *Lawfare* how the controls ratcheted up efforts to control trade in software that used zero-day vulnerabilities. BIS said the controls would require export licenses for commercially available penetration testing products and that potentially any exploit sent abroad or even to a national of a foreign country would require a license. BIS published an extensive FAQ that attempted to clarify how the controls affected security research involving exploits. The EFF criticized the FAQ for creating more confusion than clarity on the controls' scope. In addition, the Commerce Department revoked exemptions for commercially available software products that would have applied to many of the newly controlled security products. It also failed to provide license exceptions for security research. The security industry identified all of these actions as harmful to business and research activities.

But the cybersecurity industry also had problems with the substance of the Wassenaar language itself. Symantec, FireEye and other security software vendors said the intrusion software definition was too broad and it encompassed legitimate products like endpoint security systems and other tools that "hook" into a system to modify its code. They said further that the controls would also make it much more difficult for security research and vulnerability information sharing. The control on "technology for the development" of intrusion software would have covered many essential tools for the security research community such as exploit proofs-of-concept and automated vulnerability generators.
In response to a deluge of comments opposing the rule, the Commerce Department withdrew the proposal. After escalating criticism and a dressing down from the House oversight committee, the Commerce Department convened with its interagency partners to revise the U.S. approach. In March 2016, Commerce Secretary Penny Pritzker said in a letter that the United States would attempt to remove the intrusion software controls at that year's Wassenaar meetings. In December 2016, the U.S. negotiating team (with added technical experts from the cybersecurity community) failed to convince the other

40 Wassenaar members to agree on narrower language. Bipartisan groups of lawmakers in the House and Senate urged the Trump administration to continue the push to alter the Wassenaar language at the 2017 meeting.

2017 Revisions to the Wassenaar Controls

At the December 2017 Wassenaar meeting, the members agreed on a set of changes to the intrusion software controls. It received limited media coverage. The revised control list included several additions and alterations that Katie Moussouris, a security professional and technical adviser to the U.S. Wassenaar delegation, hailed as fixes for the problems the cyber industry had complained about. The changes are:

- Replacing language that controlled software "specially designed" to operate or communicate with intrusion software with the terms "software specially designed for command and control" of intrusion software.
- Adding an exception for software that carries out updates authorized by the owner or operator of the system.
- Adding exemptions for controls on technology either involved in the development of intrusion software or the development of software that operates, controls or delivers intrusion software. These exemptions said the controls do not apply for vulnerability disclosure or cyber incident response activities. The list defines vulnerability disclosure and cyber incident response as processes for sharing information about vulnerabilities and cyber incidents but does not explain how these exemptions apply to specific categories of items.
- Adding a clarifying note saying that the above-described exemptions "do not diminish national authorities' rights to ascertain compliance" with existing controls.

The alterations appear to address the some of the concerns associated with the Wassenaar language, notably the concerns from the security research community about vulnerability information sharing. But it is not yet clear how the new language mitigates concerns about the broad definition of intrusion software that may encompass legitimate security tools not used for "vulnerability disclosure" or "cyber incident response." Rob Joyce, White House cybersecurity coordinator, has praised the changes, as has Rep. Jim Langevin, a leading congressional voice on this issue.

What Comes Next?

The path forward is not clear. The Commerce Department could use the new language to craft a new proposed rule, to be followed by yet another public comment period. It would have to decide whether to add more exceptions and how to define how the new exemptions apply. Alternatively, Commerce could delay implementing the revised list and wait for the next meeting's negotiations. In that scenario, Commerce could push for the U.S. delegation to demand more substantial changes to Wassenaar's definition of intrusion software. But the human rights and internet freedom communities united with industry to oppose the 2015 proposed rule, and it is not clear whether the new changes will satisfy their concerns.

The Wassenaar changes could cause confusion for other countries as well. The EU has had intrusion software on its export control list since 2015. But as revelations that European companies sold surveillance toolkits to Middle Eastern dictators continued, the EU has sought to revamp its dual-use export control legislation in the interest of human rights. Separately, Israel (which is not a Wassenaar member but has a domestic law that adopts all Wassenaar controls automatically) attempted to more rigorously define intrusion software in early 2016. Doron Hindin detailed this effort for *Lawfare*. But a few months later, Israel shifted its policy on export controls, substantially reducing the scope and strength of

> the license requirements. It is unclear how the newest Wassenaar shift will play into both the EU export control reform initiative and the liberalized Israeli approach.

Is the problem solved? What further steps would you take?

D. What about the Trade in Personnel?

Recent events in the United States have underscored that the question of exporting hacking capability is not just a question of the trade in malware and vulns. It's also a trade in services—especially the services of trained personnel. Consider the following piece I wrote for Lawfare in February 2019:

> It's been known since 2012 that a Baltimore-based company called Cyber Point had a contract with the United Arab Emirates (UAE) to assist its newly-established signals intelligence agency (then called the National Electronic Security Authority) with "advice on cyberdefense and policy," as Ellen Nakashima reported at the time for the Washington Post. Later, there were suggestions that Cyber Point might be involved in helping the UAE service acquire malware that the UAE used to support surveillance activities that included monitoring of political opponents. And now, Reuters has a remarkable piece from Chris Bing and Joel Schactman, published last week, that goes deeper and raises important questions about the role of U.S. citizens in working for foreign intelligence agencies.
>
> The Reuters report explains that Cyber Point hired a group of ex-NSA employees to work in the UAE in support of the UAE signals intelligence service, under the name of "Project Raven." Later, the Project Raven team was transferred in some fashion from the Cyber Point contract to a contract with the UAE-based firm DarkMatter. Along the way, the Americans came to appreciate that their efforts at times did indeed include surveillance of political opponents of UAE authorities, and further that the UAE service at times targeted Americans despite assurances that this would not occur (or at least that the operations Project Raven in particular conducted or supported would not be directed at Americans). They probably should not have been surprised by any of that. But be that as it may, the story understandably has excited concern that the United States lacks a sufficient policy-and-law framework to regulate situations of this kind.
>
> What policy concerns, precisely, does this story illustrate? It's important to be clear on that question before asking whether the U.S. has an adequate legal architecture for addressing scenarios like Project Raven.
>
> I see at least four distinct areas of concern implicated by the Project Raven story, each implicating existing or potential legal architectures in different ways:
>
> 1. Whether Americans in general should ever serve in a foreign intelligence agency;
> 2. Whether former intelligence community employees in particular should not do this;
> 3. Whether the United States adequately protects classified information;
> 4. Whether the United States adequately protect the privacy of Americans from foreign surveillance.
>
> Let's take those in sequence.
>
> **1. Whether Americans in general should ever serve in a foreign intelligence agency**

The question here does not concern serving secretly as an asset for foreign governments, which U.S. law obviously should (and does) discourage. Instead, the question is whether it is inherently bad for U.S. persons (American citizens and lawful permanent residents) to have overt work relationships with foreign intelligence agencies. Some might argue that all such service is inherently unfriendly to the U.S. government at least at some level, with some contexts plainly much more so than others. Might some form of ban, with exceptions, therefore be desirable? And if so, does the U.S. not already have something like this?

The case for having some form of ban begins with the idea that there is inherent tension between owing a duty of loyalty to the United States and providing services to a foreign government in direct relation to that government's pursuit of its own security and foreign policy goals, which in some cases may be inimical to the security and foreign policy goals of the United States. There is also the possibility that the actions of U.S. persons abroad, even when working under the direction and control of a foreign government, could be held against the United States, fairly or not. On the flip side, there are of course circumstances in which interests align enough to make such service with a foreign intelligence agency beneficial for the United States. The balance of equities, then, is highly context-dependent.

This suggests that the optimal legal framework would be one that allows such service only where there has been some degree of vetting on the front end, and some degree of ongoing monitoring on the back end—that is, a licensing model, as opposed to a flat ban.

As it turns out, the U.S. already has a version of this. But as the Project Raven story itself illustrates, the current system seems incomplete.

The existing licensing model was created primarily to address military-relevant materials and services. Under the Arms Export Control Act, the executive branch is authorized to prohibit the unlicensed export of "defense articles" and "defense services." The meaning of those broad categories is explicated at *great* length and detail through the International Traffic in Arms Regulations (ITAR) and the U.S. Munitions List (USML). Things that fall within them cannot be exported without a license from the State Department's Director of Defense Trade Controls (DDTC). DDTC can—and does—impose conditions on the licenses it grants, including, where appropriate, limitations designed to prevent U.S.-provided articles and services from being used by foreign recipients in ways that violate human rights or that otherwise harm U.S. interests.

Does this apply to things or services provided to foreign intelligence services? The question is a tricky one, at least on paper, for the regulations are dense and their references to intelligence activities are sparse (and not tailored to clearly address grey zone activities in the cyber domain). That said, the Project Raven/Cyber Point episode itself makes clear that DDTC does view at least some such activity as coming within the licensing system; Cyber Point applied for and received a license, after all.

The more important questions—the ones the Project Raven story raises in relation to this particular policy concern—are whether the licensing system is sufficiently probing on the front end when the license is being granted, and whether there is adequate back-end monitoring for compliance with license conditions. The nuances in the Project Raven story, so far can be gleaned from the Reuters account, are tricky on both points.

Reuters writes that Project Raven transitioned out from under Cyber Point (and its license) at some point, moving to UAE-based DarkMatter. Conduct subsequent to that shift probably cannot be laid at Cyber Point's door, and thus would not violate the Cyber Point

license. On that view, one cannot fault the front-end screening by DDTC, unless it turns out that the Project Raven activities violated the license terms all along. On the other hand, that same logic compels the conclusion that the U.S. persons who remained working for the UAE at that point no longer had the benefit of a DDTC license, yet they continued providing "defense services" to the UAE. This suggest a possible enforcement gap in the licensing scheme, though perhaps time will yield an enforcement gap now that the story has become public.

At any rate, one potentially-attractive policy response to the whole episode would be to increase the resources devoted to both front-end screening and in-progress monitoring for DDTC licenses. And perhaps the statute could be amended to support that in-progress monitoring by imposing increased requirements for periodic compliance-reporting.

Before moving on, it's worth noting how things might look if the country decided licensing was not the right way to go and, instead, wanted to move to a flat ban. That idea has a close parallel in the context of foreign *military* service. In that setting, American criminal law has long made it a crime to enlist or otherwise enter into the service of a foreign military (unless that foreign military is at war with a state against which the U.S. too is at war), and separately it has also long been a crime to take a foreign government's commission to serve in war against a party with which the United States is at peace. To be sure, there are ample historical examples of the government looking the other way; enforcement has not been anything like uniform or rigid over time. But still, there it is.
Should the U.S. treat foreign *intelligence* service the same way as foreign *military* service, with a flat ban? This would certainly help minimize Project Raven scenarios in which U.S. persons end up contributing to undesirable foreign intelligence activities. But perhaps it would overcorrect to too great an extent, for it also would preclude U.S. persons from contributing to *desirable* activities, such as the sort of counterterrorism functions that the Project Raven personnel originally thought would be their focus. A well-tailored and well-resourced licensing system seems to me to be the better alternative.

2. Whether former intelligence community employees in particular should not do this
Even if Americans in general should have the option of foreign intelligence service, subject to proper licensing, one might argue that former intelligence community employees are a special case for whom no such license should be granted.

The case for a flat-ban for former intelligence community employees might go something like this: First, as the Project Raven story suggests, the fact that an American working for a foreign service once was an intelligence community employee considerably enhances the extent to which the actions of the foreign entity may come to reflect back on the United States, even if unfairly so. Second, much like the analogous scenario in which bans or at least delays are imposed before former public officials can engage in lobbying related to their former jobs, the U.S. should worry about both the appearance of impropriety (which undermines public regard for the entity in question, and thus inhibits that entity's ability to pursue its mission) and the possibility that people in public service might be influenced in their decisions by future employment prospects. Third, officials might worry (as noted below) about the opportunities for compromise of U.S. personnel that are created by working with a foreign intelligence service. (Indeed, it is not hard to see how the UAE service could have taken advantage of the growing murkiness of the Project Raven activities so as to create leverage over at least some of those former NSA employees).

The case against such a flat ban, in contrast, builds from the premise developed above: there are some circumstances in which it is in the U.S. national interest to improve the

efficacy of foreign intelligence services by allowing Americans to work with them. If that's the case, it stands to reason that the persons most able to provide that boost, in at least some cases, will be former intelligence community employees. On that view, the licensing system to must function well enough to police against undue risk of the kinds just noted in the preceding paragraph.

3. Whether the U.S. adequately protects classified information

The preceding discussion draws attention to a related, but distinct, concern: Does the Project Raven scenario highlight an undue risk that classified information involving cyber capabilities will leak to a foreign service as a result of former intelligence community personnel serving with a foreign agency?

Yes and no. The good news is that the U.S. doesn't lack for relevant criminal laws in this area. Though there always is risk that a former employee will choose to violate those laws or be compromised into doing so, there's not much cause for trying to expand or strengthen the legal guardrails when it comes to obvious concerns such as exposure of specific tools like custom malware, access to staging servers, specialized physical equipment, etc. And the Project Raven story does not suggest otherwise.

But on the other hand, there also is great value in the sheer practical knowledge—the tactics, techniques and procedures (TTPs)—for which the foreign agency is hiring these former intelligence community personnel in the first place. Whether and when that know-how is itself classified information can be tricky, to say the least. Knowledge concerning a specific software vulnerability—of a zero day—might readily count, but TTPs involving best practices for detection avoidance and lateral movement within a system might present a murkier case. And the less clear things are, the more likely it is the former employee will simply use the TTP in the new job, with the new employer and new colleagues gaining that know-how along the way even if the U.S. person never intended to "train" them in any formal sense.

It seems to me the TTP-sharing issue is intractable to a certain extent. If one takes a sweeping approach to classifying know-how and then ensuring intelligence community employees understand there will be a strict approach to enforcement, this runs the risk of effectively prohibiting those employees from going on to do related work for *anyone* outside the government, not just foreign intelligence agencies. This in turn would make the already-serious challenge of recruiting talented hackers and defenders much more so.

4. Whether U.S. law adequately protects the privacy of Americans from surveillance by Americans working for a foreign government

Perhaps the most striking element of the Project Raven story is the reference to UAE surveillance of Americans. My read of the story was that the Americans were not themselves engaging in this activity. If that's right, then all that's at issue is the fact that foreign intelligence services spied on U.S. citizens—an important thing to know, but not something that warrants some innovation in the U.S. legal architecture). But what if that's not right? That is, what if some of the Americans associated with Project Raven were engaged in surveillance of their fellow citizens, on behalf of a foreign government?

The short and complete answer is that they would be in serious legal jeopardy, for neither their license (which surely excludes such activity anyway) nor the cloak of UAE domestic law would do anything to make such activity legal from the perspective of U.S. law. Various U.S. laws—the Wiretap Act, for example, and the Computer Fraud and Abuse Act—might come into play.

> Indeed, the Reuters report notes that the FBI has taken a keen interest in Project Raven participants. Perhaps this is one of the reasons why? Time will tell. For now, it is enough to say that this is not an area where the legal framework seems lacking.

Does the United States need new laws?

22. Government Hacking: Law Enforcement

Up to this point we have been considering the potential desirability and legality of allowing private-sector entities to hack in certain circumstances. But what about the government itself?

A. What Policy Goals Might Government Hacking Serve?

Let's begin by identifying some of the overarching policy goals that might be advanced by allowing certain government agencies latitude to hack. Here is a non-exhaustive list of possibilities:

1. **Law Enforcement:** hack to gather evidence of crime
2. **Espionage:** hack to steal foreign-intelligence information
3. **Armed Conflict:** hack in direct relation to an armed conflict
4. **"Gray zone" Competition:** hack to cause effects on foreign adversaries short of war
5. **Nonconsensual Defense:** hack to patch or defend others involuntarily

In each case listed above, the general idea is that in some circumstances it might be efficient for the government to pursue certain policy goals via hacking (as opposed to having to rely on the various other methods otherwise available to the government).

Can you imagine a hypothetical example, for each category, in which hacking might be the most efficient way for the government to proceed?

Knowing that hacking might sometimes be convenient for the pursuit of important policy goals only starts the analysis, of course. Government hacking might entail offsetting costs, too.

For each of your hypothetical examples, can you identify at least one cost or risk that use of a hacking approach (in contrast to some other approach) might entail?

Insofar as the costs or risks seem to outweigh the benefits of government hacking in at least some of these circumstances, the question then becomes whether we can reduce those costs and risks to a tolerable level by adopting various rules and procedures intended to facilitate government hacking while safeguarding against anticipated harms. Bear that in mind through the remainder of the readings, as we survey the circumstances in which U.S. government institutions currently have (or might one day have) authority to hack.

We will proceed in the same order as the list above, beginning with government hacking to gather evidence of crime.

B. Law-Enforcement Hacking

In the abstract, criminal investigation can be understood as a sometimes-complex process of gathering information and then applying some combination of expertise, reason, and analytic technique to gain insight from it. In today's world, that information-gathering function often entails seeking access to data that is "at rest" (running the gamut from a business's records to communications that are stored on some server or device) or that is "in motion" at the time of acquisition (think, for example, of wiretapping to gain access to the content of an email).

The government's legitimate criminal law investigative interests are in tension, of course, with the privacy interests of the persons and entities whose data and communications might be at issue. Not surprisingly, the United States has developed a fairly elaborate legal framework to mediate that tension.

1. A Thumbnail Sketch of the Legal Framework for Criminal Investigations; or, Why Hack?

The Fourth Amendment to the Constitution of the United States is a foundational part of that framework. It has two key parts. First, it requires that all government actions constituting a "search" be "reasonable." In practical terms, that means in *most* cases the government first must obtain a "warrant" from a judge before conducting a "search." Second, the Fourth Amendment also specifies that judges may not issue a warrant unless the government has provided a sworn statement that is sufficient to establish "probable cause" that the search will yield evidence of (or fruits from) a crime, and that "particularly" describes the "place to be searched" and the "things to be seized."

But not every search must be supported by a warrant to be count as reasonable. The Supreme Court over time has recognized various exceptions and other limits to the reach of that rule. Most notable for our purposes, the Supreme Court in the 1970s adopted the "third-party doctrine." Under that rule, a person has no reasonable expectation of privacy in information that is in the custody of a third party (such as your bank or your phone company), and thus the Fourth Amendment has no application should the government ask such a third party to turn over that information to it.

This rule has been controversial from the start for obvious reasons, and in 2018, the Supreme Court caused a commotion in *Carpenter v. United States* by declining to apply the third-party doctrine to a circumstance in which the government obtained from a telecommunications company an entire month's worth of cell-site location data that showed a suspect's movements in a particularly comprehensive way. For the time being, however, the third-party doctrine remains the rule for *most* investigation scenarios (indeed, the Supreme Court in *Carpenter* itself went out of its way to suggest that traditional investigative scenarios remain subject to this rule).

Against the backdrop of this limitation on the relevance of the Fourth Amendment, Congress over time has enacted a variety of statutes that both limit and facilitate government requests to third parties for production of information (including electronic information at rest and in motion). It is beyond the scope of this course to explore those rules in detail, but for present purposes it is enough to know that criminal investigators routinely seek production of information from third parties via instruments that require less of an evidentiary showing than does a warrant and that do not always require ex ante judicial involving, including "grand jury subpoenas" and orders issued by courts under the 1986 Stored Communications Act. What's more, Congress even imposed requirements intended to help overcome potential technical obstacles that might arise in such cases. For example, the Communications Assistance to Law Enforcement Act (CALEA) requires telecommunication companies to "ensure" that their systems are "capable of" compliance with court-ordered wiretaps and the like. And the Wiretap Act further provides that:

> An order authorizing the interception of a . . . communication under this chapter shall, upon request of the applicant, direct that a provider of wire or electronic communication service, landlord, custodian **or other person** shall furnish the applicant forthwith all information, facilities, and **technical assistance necessary to accomplish the interception** unobtrusively and with a minimum of interference with the services that such service provider, landlord, custodian, or person is according the person whose communications are to be intercepted. (emphasis added).

In light of all this, it might seem that the government would have little occasion to engage in hacking. Why sneak in the back door, after all, if the front door is held open for you by the force of law? The first and foremost answer is that the front door is not always open as a *practical* matter, even if it is open as a *legal* matter.

How might that be the case? Without looking ahead to the readings that follow, can you identify any scenarios that might fit this description?

2. Lawful Hacking

To understand how FBI eventually came to consider hacking as an option to overcome practical impediments to the execution of lawful investigative authorities, some historical context is in order.

In the 1990s, FBI developed a "traffic sniffer" tool intended to be used in support of lawfully authorized surveillance. The tool—unwisely named "Carnivore"—provided a capacity to filter packets transiting the network of an Internet Service Provider, flagging those that were to or from a lawfully-approved target for surveillance. All of this would occur in cooperation with the ISP, so it was not an example of FBI hacking. Rather, it was an example of the FBI developing a capacity to accomplish something in circumstances in which the ISPs themselves either could not or would not develop that same capacity (ISPs, notably, do not count as telecommunication companies for purposes of CALEA).

When the public became aware of Carnivore, it proved exceptionally controversial. Some objected to the possibility that it might result in collection of much more than the target's communications, while others advanced more sweeping objections. Eventually, though, the issue would become moot. As Kim Zetter explains in this article from Wired, the increasing ubiquity of encryption began making it impossible (or at least much harder) for FBI to derive readable plain-text from the packets that might be flagged via Carnivore or similar technologies.

How to overcome that obstacle? Some of the options entailed obvious downsides, not to mention posed massive political obstacles. One might try to ban encryption altogether, for example, or to compel providers of encrypted communications services to maintain a capacity to decrypt upon court order. But another option was to shift focus away from intercepting communications in transit (at which point encryption was an issue), and instead work to gain surreptitious access to the sender or receiver's computer and thus see the pre-encryption or post-decryption messages in plain text after all. This was a much taller order and less scalable solution, to be sure, but it proved tempting enough to try.

Zetter's article highlights a 1999 investigation of an organized crime figure as "the first criminal suspect known to be targeted by a government keystroke logger"—*i.e.*, malware that would create a record of all the typing that occurs on a computer. The stakes were high in that case, for the agents involved had to break into the suspect's home (twice!) in order to install the logger.

They obtained a warrant, however, and managed the trick. And soon break-ins would no longer be needed, as FBI developed malware that could be delivered remotely. The age of "lawful hacking" was underway, eventually expanding to encompass a variety of investigative capacities beyond keylogging.

Consider the following questions:

- Is it correct to call this "lawful" hacking? Why or why not?
- Can you articulate how the ubiquity of encryption drives law enforcement interest in hacking?
- Can you relate this law enforcement interest back to the topic of our last class: the insecurity industry?

3. Watering Holes and Network Investigative Techniques

Notice that our examples, up to this point, involved known suspects. What happens when the *crime* is apparent, but the *suspect* is not? Consider the issues that arose in *United States v. Henderson*, a decision by the United States Court of Appeals for the Ninth Circuit upholding the legality of a law-enforcement hacking tactic used in connection with a child pornography investigation:

> In 2014, the Federal Bureau of Investigation ("FBI") began investigating the internet website upf45jv3bziuctml.onion, "Playpen," which was used to send and to receive child pornography. Playpen operated on an anonymous network known as "The Onion Router" or "Tor". To use Tor, the user must download and install the network software on his computer. Tor then allows the user to visit any website without revealing the IP address, geographic location, or other identifying information of the user's computer by using a network of relay computers. Tor also allows users to access "hidden services," which are websites that are accessible only through the Tor network and are not accessible publicly. A hidden-service website hosted on the Tor network does not reveal its location; a Tor user can access the hidden-service website without knowing the location of its server and without its knowing the user's location. Playpen operated as a hidden-service website and required users to log in with a username and password to access its discussion forums, private messaging services, and images of child pornography.
>
> After determining that Playpen was hosted on servers located in Lenoir, North Carolina, the FBI obtained and executed a valid search warrant in the Western District of North Carolina in January 2015, and seized the Playpen servers. The FBI removed the servers to its facility in Newington, Virginia. Because Tor conceals its users' locations and IP addresses, additional investigation was required to identify Playpen users. The FBI then operated the Playpen website from a government-controlled server in Newington in the Eastern District of Virginia, from which it obtained a valid court order authorizing it to intercept electronic communications sent and received by the site's administrators and users.
>
> The FBI later obtained a warrant from a United States magistrate judge in the Eastern District of Virginia … authorizing searches for thirty days using what is known as a Network Investigative Technique ("NIT"). Specifically, such "NIT warrant" authorized the search of all "activating" computers—that is, those of any website visitor, wherever located, who logged into Playpen with a username and password. The NIT technology is computer code consisting of a set of instructions. When a person logged into the Playpen site, the NIT caused instructions to be sent to his computer, which in turn caused the computer to

respond to the government-controlled server with seven pieces of identifying information, including its IP address. The NIT mechanism allowed the FBI, while controlling the website from within the Eastern District of Virginia, to discover identifying information about activating computers, even though Playpen operated on the Tor network.

On March 1, 2015, a person logged into Playpen under the username "askjeff." The NIT instructions were sent to askjeff's computer, which revealed its IP address through its response to the government-controlled server. The computer response also revealed that askjeff had been actively logged into Playpen for more than thirty-two hours since September 2014 and had accessed child pornography.

The FBI traced the IP address to an internet service provider ("ISP"), Comcast Corporation, which was served with an administrative subpoena requesting information about the user assigned to the IP address. The IP address turned out to be associated with a computer at the San Mateo, California, home of Bryan Henderson's grandmother, with whom Henderson lived. A local federal magistrate judge in the Northern District of California issued a warrant to search the home, where the FBI then discovered thousands of images and hundreds of videos depicting child pornography on Henderson's computer and hard drives. [Henderson was convicted, and later appealed the trial court's denial of his motion to suppress the NIT-derived evidence.]

. . . Henderson challenges only the warrant issued by the Eastern District of Virginia on February 20, 2015, authorizing the use of the NIT. . . . Henderson argues that . . . the NIT warrant was issued in violation of Federal Rule of Criminal Procedure 41(b), which authorizes magistrate judges to issue warrants subject to certain requirements. . . .

Henderson urges that no provision within Rule 41(b) authorizes a magistrate judge to issue the NIT warrant to search computers located outside of her district. In general, Rule 41(b) permits "a magistrate judge with authority in the district . . . to issue a warrant to search for and seize a person or property located within the district."

. . . The government concedes that a "search" occurred when the NIT was deployed to users' computers and returned their identifying information. As two of our sister circuits have before us, we agree. . . . However, the government counters that the NIT warrant was nonetheless authorized under Rule 41(b)(4)'s specific provision for tracking devices, which permits "a magistrate judge with authority in the district . . . to issue a warrant to install within the district a tracking device . . . to track the movement of a person or property located within the district, outside the district, or both."

Rule 41 defines a "tracking device" as "an electronic or mechanical device which permits the tracking of the movement of a person or object." The government contends that Henderson's computer made a "virtual trip" to the government server in the Eastern District of Virginia when he logged into the Playpen website. According to the government, his computer then "brought" the NIT instructions, along with the usual Playpen website content, back with it from the government server to his computer's physical location in California. The NIT instructions then caused identifying location information to be transmitted back to the government, just like a beeper or other tracking device would.

We are not persuaded by the government's assertions. The NIT instructions did not actually "track the movement of a person or property," as required by the tracking-device provision. Rather, the NIT mechanism was simply a set of computer instructions that forced activating computers, regardless of their location, to send certain information to the government-controlled server in Virginia. Users' computers did not physically travel to

> Virginia, and the information they relayed did not reveal the physical location of any person or property, unlike a beeper attached to a vehicle. The "seized information (mainly the IP address) assisted the FBI in identifying a user, [but] it provided no information as to the computer's or user's precise and contemporary physical location."
> . . .
>
> Interestingly, Rule 41(b) was amended on December 1, 2016—after the issuance of the NIT warrant here—to authorize magistrate judges to issue warrants to search computers located outside their district if "the district where the media or information is located has been concealed through technological means." As our sister circuits have recognized, such amendment plainly seems to "authorize[] warrants such as the NIT warrant here." . . . The fact that Rule 41 was amended to authorize specifically these sorts of warrants further supports the notion that Rule 41(b) did not previously do so.
>
> . . . But does a warrant issued in violation of Rule 41(b) compel suppression of evidence? Not necessarily. Only certain Rule 41 violations justify suppression. The suppression of evidence is "a judicially created remedy designed to safeguard Fourth Amendment rights generally through its deterrent effect, rather than a personal constitutional right of the party aggrieved." . . . [The Ninth Circuit went on to conclude that the suppression was not warranted here because the officers relying on the warrant did so out of good-faith belief in its validity, and because there was no forward-looking deterrence justification for suppression given the intervening amendment to Rule 41 noted above.]

Consider the following questions:

- How would you explain, in plain terms, what the FBI did here under the label "NIT"?
- If a private person did the same thing, would it violate the CFAA?
- Can you identify any plausible alternative way for the FBI to identify the persons accessing Playpen other than the NIT approach they used?
- Can you identify any undesirable policy consequences to allowing the FBI to use this method? Don't forget: they did get a warrant.
- Change the facts: What legal consequence(s) if FBI had not obtained a warrant?
- Does it matter to your analysis that Playpen preexisted? What if FBI created it from scratch, as a "watering hole" tactic?
- Revisit the topic of vulnerability disclosure, which arose with the last reading. Does our foray into NITs and lawful hacking change your views on that topic?

4. Hacking to Get Data-at-Rest (and the Going-Dark Connection)

Note: The NIT language is often used as a shorthand for lawful hacking scenarios, but not all law enforcement hacking is network-focused. Sometimes, for example, the idea is to overcome the security of an encrypted device in order to access data-at-rest (or, as in the mafia case mentioned earlier, to install the equivalent of a keylogger in order to catch outbound and inbound traffic pre- and post-encryption). The "Apple v. FBI" controversy looms very large under that heading. As you probably are aware, that controversy is closely associated with a larger policy debate sometimes called the "Going Dark" debate.

The idea with that label is that law-enforcement agencies are losing the practical ability to end up with readable plain-text even when they have warrants or other lawful bases for accessing data in motion or at rest, because of the combination of two trends: (1) the growing ubiquity of

encryption as the default for the varied communications platforms people use these days, and (2) circumstances in which the company that created the relevant communication platform or device does not have—and is not willing or perhaps even able to create—a workaround. Thus, a series of federal law-enforcement officials have argued that Congress should enact a new and broader CALEA-style statute, this time focused on requiring companies to have and preserve the ability to decrypt traffic or otherwise assist law enforcement in executing warrants in this situation. Opponents argue that there is naught that can or should be done about this, asserting that the net gain for security overall outweighs lost investigative benefits, and that law enforcement in the meantime may be enjoying a "golden age" of investigative benefits in relation to non-content metadata that remains available.

A proper dive into the Going Dark debate is beyond the scope of our course, but if you want to spend more time on your own studying these issues, consider reading the "Don't Panic" report issued by the Berkman Center at Harvard and this essay from Susan Hennessey for Brookings, and listening to this special episode of Pat Gray's Risky Business podcast (part of a special series sponsored by the Hewlett Foundation) in which Pat interviews former FBI General Counsel Jim Baker on the subject.

23. Government Hacking: Espionage

Now let's turn our attention away from law-enforcement hacking, focusing instead on hacking carried out by the government's Intelligence Community for purposes of espionage—that is, "intelligence collection" conducted not to make criminal cases but rather to inform the decisions and plans of our nation's leaders (something we first described way back in class #6 when, as part of our effort to understand the attacker's perspective, we were distinguishing among various government purposes for hacking).

Not all espionage takes the form of hacking, obviously. The traditional intelligence-collection "disciplines" include, among others, the following approaches:

1. **Human intelligence (HUMINT)**, which involves inducement of a human source, with the inducement typically involving one of more of: money, ideology, coercion, or ego;
2. **Signals intelligence (SIGINT)**, which involves interception of radio or other electromagnetic-spectrum signals (both in the form of words and otherwise);
3. **Geospatial intelligence (GEOINT)**, which involves satellite and aerial imagery as well as mapping data; and
4. **Collection of tangible things** (via break-ins or otherwise).

Some would argue that hacking conducted by an intelligence agency for purposes of espionage should be categorized as an intelligence discipline all its own. That said, it probably fits best as a subcategory of SIGINT. Whatever the label, though, there's no question that it has become a central part of the espionage toolkit for every government that engages in spying.

A. Which U.S. Government Institutions Perform this Function?

The key U.S. government agencies for purposes of this function are, not surprisingly, the NSA and the CIA.

Both are part of the broader U.S. government "Intelligence Community," a label often abbreviated as "the IC" (pronounced eye-see). The IC refers to the full array of federal agencies that engage in intelligence activities. There are *seventeen* component parts currently. Eight of them are part of the Defense Department (including, most notably for our purposes, the National Security Agency), and these all fall within the budget, policy, and personnel domain of the Secretary of Defense. Seven others are part of other full-fledged government departments, including the Departments of Justice, State, Treasury, and Homeland Security. Only two stand as independent agencies: the CIA, and the Office of the Director of National Intelligence (ODNI).

ODNI was created in 2004 with the goal of providing IC-wide coordination and services, and thus the Director of National Intelligence ("DNI") to a *limited* extent functions as the head of the IC. The DNI's control over other IC institutions is quite limited both formally and functionally, however. If one really wants to describe how authority over the IC is distributed, frankly, it probably is most accurate to say that the DNI combines with the Director of the CIA and especially the Secretary of Defense (or, on a day-in, day-out basis, the Under Secretary of Defense for Intelligence) to form a sort of informal triumvirate of senior-most intelligence officials.

From a "Unit I" (defensive) cybersecurity perspective, all seventeen components of the IC are relevant to this class because collection and analysis of cybersecurity threats looms so large to the defensive mission. But most IC components do not have an *offensive* cyber-domain role in the sense of conducting hacking for espionage purposes. Both NSA and CIA, however, very much have such a role.

We last checked in on NSA in reading #16, as we took note of NSA's cybersecurity directorate and its efforts to boost cybersecurity defense in the United States in certain situations. That is an important function for NSA, but make no mistake: NSA's core function is SIGINT collection and analysis, for purposes of both foreign-intelligence gathering in general and also collection of intelligence for combat-support purposes (don't forget, NSA remains part of the Defense Department). And toward these ends (as well as its defensive mission), it has long maintained world-class capacities for hacking (including extraordinary in-house capabilities for identifying vulns and crafting exploits). Note that the Director of the NSA also serves, simultaneously, as the commander of United States Cyber Command ("CYBERCOM), which we will study in our next session; this arrangement is called "the dual hat."

The Central Intelligence Agency ("CIA") is most famous for its HUMINT-collection activities and its analytical prowess, but it also has its own in-house capacities for developing hacking tools and implementing them in the field.

B. What Legal Framework Regulates this Activity?

Over the past five decades, the United States has developed a complex legal framework relating to espionage activities. Though it does not include provisions that are specific to espionage conducted via hacking, its general provision applies in that context the same as with any other espionage activities.

A full study of that framework is beyond the scope of this course (see my eCasebook Chesney on the Law of the Intelligence Community for more). There are some highlights, however, you should know.

Like most legal frameworks pertaining to government activity, the legal architecture for espionage addresses three types of questions:

1. Which agencies have affirmative authority to engage in certain kinds of activity?

2. What process must be followed in order for an otherwise-authorized agency to use its authority?
3. And what substantive limits does the law place on the resulting activity?

Let's have a quick look at each of these.

1. Authority to Act

There is no serious dispute about the affirmative authority of certain IC members to engage in collection. Take the CIA: Congress has expressly authorized it to "collect intelligence through human sources and by other appropriate means," 50 USC 3036(d)(1), and has appropriated considerable sums for this purpose since the mid-20th century. Even if this were not the case, moreover, the executive branch undoubtedly would assert inherent authority to engage in foreign-intelligence collection under Article II of the Constitution, citing the president's duties in relation to both foreign affairs and national defense (note that this is quite different—and much less controversial—than asserting inherent authority also to override a statutory constraint; it is simply a claim that Congress's affirmative permission is not needed in order to engage in foreign-intelligence collection (though Congress's *money* may well be needed!)).

As for NSA? The situation there is somewhat different, for there is no comparably clear statutory statement spelling out NSA's various missions. There is a comparable history of Congressional funding and oversight, however, not to mention a deep history of presidents implicitly asserting Article II authority to order the military to conduct SIGINT collection (for NSA is the direct institutional successor to Army and Navy entities that independently performed SIGINT functions in the first half of the 20th century).

The interesting questions about CIA and NSA collection, as we shall see, tend to concern not authorization as such, but rather the rules of process and substantive constraints spelled out below.

2. Procedural Requirements

Congress has passed a number of statutes regulating the *process* of engaging in espionage. Some of these rules control the *ex ante* process of deciding to engage in some particular activity in the first place. For example, must the executive branch obtain approval from a judge, or must some particular executive branch official approve? And some involve *ex post* oversight rules involving requirements to keep Congress informed.

On the front end (prior to a collection operation), the interesting issue is whether and when the government must obtain *judicial* approval. For our narrow purposes, the answer is no; Congress has never subjected NSA or CIA hacking activity directed at foreign targets outside the United States to any such requirement. If we were instead exploring other forms of collecting foreign-intelligence information, such as compelling a U.S. company to cooperate in wiretapping, the answer would be quite different (and complex).

Congress also has not mandated any particular decision-making process be followed within the executive branch for purposes of deciding to engage in espionage (as we will see in the next reading, the situation is different with respect to the separate category we call "covert action").

What Congress *has* done, though, is imposed a requirement on the DNI and also the head of any relevant agency (like NSA or CIA) to keep the House and Senate Intelligence Committees "fully and currently informed" of collection activities, "including any significant anticipated" activities and significant "failures" too. So it says in 50 USC 3092, which sets forth this rule for all intelligence activities aside from covert action (since covert action is subject to its own oversight rules, as we will see in the next reading). And 50 USC 3091 imposes a similar obligation on the President.

Consider the following questions:

- Can you explain how the mere existence of a requirement to report to Congress on collection activities might have a beneficial effect in terms of deterring bad ideas?
- Can you think of any costs to having such reporting requirements (and do you need to know more about the granularity of that reporting before you can answer that question)?

3. Substantive Constraints

Congress also can regulate intelligence activities by placing certain actions off-limits altogether. That is, Congress could specify certain things that NSA, CIA, and other intelligence agencies simply may not do when conducting espionage. It could even impose substantive constrictions that would be specific to the context of espionage conducted via hacking. For example, Congress could prohibit any steps involving the hacking of an Industrial Control System associated with a foreign country's electrical grid. Thus far, however, Congress has taken no such steps.

Consider the following questions:

- Should Congress take that exact step?
- If it did, do you think Russia, Iran, North Korea, or China would reciprocate?
- If any of them did, do you think that such a law would be obeyed?

We should pause here to note that the executive branch can impose rules on *itself*, even when Congress has not seen fit to do so. Indeed, amidst the controversies unleashed by the Snowden revelations that began in the summer of 2013, President Obama in Presidential Policy Directive 28 ("PPD-28") articulated several such constraints for SIGINT collection, and they remain on the books at this time. They include:

> **Section 1(b)** – The United States shall not collect signals intelligence for the purpose of suppressing or burdening criticism or dissent, or for disadvantaging persons based on their ethnicity, race, gender, sexual orientation, or religion. Signals intelligence shall be collected exclusively where there is a foreign intelligence or counterintelligence purpose to support national and departmental missions and not for any other purposes.
>
> **Section 1(c)** - The collection of foreign private commercial information or trade secrets is authorized only to protect the national security of the United States or its partners and allies. It is not an authorized foreign intelligence or counterintelligence purpose to collect such information to afford a competitive advantage [*FN*: "Certain economic purposes, such as identifying trade or sanctions violations or government influence or direction, shall not constitute competitive advantage."] to U.S. companies and U.S. business sectors commercially.

Consider the following questions:

- Is anything in Section 1(b) a serious constraint on what the United States otherwise might genuinely be inclined to do?
- What about Section 1(c), bearing in mind the explanatory footnote?
- Should any of this be changed?
- Do you think any other country that engages in serious SIGINT collection against the United States actually follows comparable rules?

PPD-28 goes on in Section 2 to discuss a new limit on the use of "bulk collection" of signals intelligence. Such collection need not involve hacking, but it might. For example, NSA in theory might have the capability to surreptitiously access the network of a foreign Internet Service Provider, and from that perch collect all traffic crossing that network, relying on post-collection querying of the bulk results in order to separate the wheat from the chaff. At any rate, Section 2 of PPD-28 articulates limits on such approaches, whether made possible by hacking or not:

> Locating new or emerging threats and other vital national security information is difficult, as such information is often hidden within the large and complex system of modern global communications. The United States must consequently collect signals intelligence in bulk in certain circumstances in order to identify these threats. Routine communications and communications of national security interest increasingly transit the same networks, however, and the collection of signals intelligence in bulk may consequently result in the collection of information about persons whose activities are not of foreign intelligence or counterintelligence value. The United States will therefore impose new limits on its use of signals intelligence collected in bulk. These limits are intended to protect the privacy and civil liberties of all persons, whatever their nationality and regardless of where they might reside.
>
> In particular, when the United States collects nonpublicly available signals intelligence in bulk, it shall use that data only for the purposes of detecting and countering:
>
> (1) espionage and other threats and activities directed by foreign powers or their intelligence services against the United States and its interests;
> (2) threats to the United States and its interests from terrorism;
> (3) threats to the United States and its interests from the development, possession, proliferation, or use of weapons of mass destruction;
> (4) cybersecurity threats;
> (5) threats to U.S. or allied Armed Forces or other U.S or allied personnel; and
> (6) transnational criminal threats, including illicit finance and sanctions evasion related to the other purposes named in this section.
>
> In no event may signals intelligence collected in bulk be used for the purpose of suppressing or burdening criticism or dissent; disadvantaging persons based on their ethnicity, race, gender, sexual orientation, or religion; affording a competitive advantage to U.S. companies and U.S. business sectors commercially; or achieving any purpose other than those identified in this section.

Consider the following questions:

- Given the six approved categories, is there anything significant left out?
- If so, is that a good thing or a bad thing?
- Do you suppose any other country that engages in SIGINT collection against the United States similarly constrains its use of a bulk-collection methods?

24. Government Hacking: Armed Conflict

In addition to hacking for purposes of criminal investigations and espionage, the U.S. government also might wish to engage in hacking for a variety of other reasons including:

- as part of an armed conflict; or
- as a means to engage in competition with foreign actors below the threshold of armed conflict (including both acknowledged and deniable actions).

We will conclude the course in this reading and the next by surveying the legal and policy issues associated with U.S. government activity (and that of other governments) in these contexts, starting with armed conflict.

A. Cyberspace operations in the context of armed conflict

Both regular armed forces and non-state armed groups increasingly depend on computers and communication networks for a growing array of functions. It follows inevitably that military organizations will seek to establish the ability to penetrate adversary systems. This certainly is true for the armed forces of the United States.

To some extent, military hacking is a topic that falls within the scope of the espionage discussion in the previous reading. Military organizations often perform intelligence collection, after all, both in connection with armed conflict (as a form of combat support) and otherwise (as part of the process of understanding the capabilities and intentions of potential adversaries). But not all hacking carried out by military entities fits that bill. Military hacking also might occur in order to achieve what we might call an "operational effect." For example, a military organization might hack a system in order to manipulate information stored there, to disrupt or destroy its functionality, or to hold it "at risk" of such an effect.

This brings us to a critical point: military organizations conduct cyberspace operations in connection with armed conflicts, yes, but not *only* in connection with armed conflicts. Knowing that a particular cyberspace operation was conducted by a military organization thus is not enough, on its own, to tell us whether we are dealing with a circumstance involving armed conflict (or to put things in more colloquial terms, it is not enough to tell us whether we are dealing with a circumstance involving war).

For the moment, we will focus on military hacking in circumstances that *do* constitute armed conflict. We will save until the next subsection the critical and recurring set of issues that arise when states use hacking (using military organizations or otherwise) to cause operational effects in the cyber domain below that threshold.

B. The role of CYBERCOM

We have discussed US Cyber Command previously in connection with our study of how the U.S. government defends its own networks (recall that JFHQ-DoDIN, the entity with overall responsibility for defending the DoDIN, is a component of CYBERCOM), and also our study of the larger defensive ecosystem and the way it can interact with international relations (for CYBERCOM at times will take public-protecting actions such as providing IOCs to VirusTotal, especially (though apparently not only) when accompanied by a name-and-shame attribution to a foreign government). But CYBERCOM's mission is by no means limited to defense. Bear that in mind as you read this thumbnail sketch of CYBERCOM's origins.

As Fred Kaplan explains in his book *Dark Territory: The Secret History of Cyber War* (available here if you are interested in going deeper), the military's effort to organize for cyber operations traces back to the late 1990s. As the Department began to appreciate how vulnerable its own networks were, it established a new office (the "Joint Task Force—Computer Network Defense," or just "JTF-CND") to coordinate defensive efforts. Kaplan writes that the

> "initial plan was to give [JTF-CND] an *offensive* role as well, a mandate to develop options for attacking an adversary's network…. [But the organizer] knew that the services wouldn't grant such powers to a small bureau with no command authority. … [Eventually, in] 2000, JTF-CND became JTF-CNO, the *O* standing for "Operations," and those operations included not just Computer Network Defense but also, explicitly, Computer Network *Attack*…. [JTF-CNO] was placed under the purview of U.S. Space Command…it was an odd place to be, but SpaceCom was the only unit that wanted the mission…[and] in any case, it was a *command*, invested with war-planning and war-fighting powers. [But key leaders] felt that the cyber missions—especially those dealing with cyber *offense*—should ultimately be brought to the Fort Meade headquarters of the NSA." (pp. 121-22)

It took many years, but that is what happened in the end. In the summer of 2009, Secretary of Defense Gates directed the creation of a new command—United States Cyber Command (CYBERCOM)—focused on both defensive and offensive functions. In order to ensure its rapid maturation, moreover, the new command would be collocated with NSA at Ft. Meade, and NSA's Director would be "dual-hatted" as the CYBERCOM commander as well. This would enable NSA to incubate CYBERCOM in terms of personnel, knowledge, technical capabilities, and so forth. Less obviously, it also would ensure a process for deconfliction of priorities should the interests of CYBERCOM in causing an operational effect in cyberspace come into conflict with the interests of NSA in collecting intelligence.

So, what exactly is CYBERCOM's role? The Defense Department's 2015 Cyber Strategy document provides a handy explanation:

> *"In 2012, DoD began to build a [Cyber Mission Force ("CMF")] to carry out DoD's cyber missions. Once fully operational, the CMF will include nearly 6,200 military, civilian, and*

contractor support personnel from across the military departments and defense components…. The Cyber Mission Force will be comprised of cyber operators organized into 133 teams, primarily aligned as follows:

> ***Cyber Protection Forces** will augment traditional defensive measures and defend priority DoD networks and systems against priority threats;*
>
> ***National Mission Forces** and their associated support teams will defend the United States and its interests against cyberattacks of significant consequence; and*
>
> ***Combat Mission Forces** and their associated support teams will support combatant commands by generating integrated cyberspace effects in support of operational plans and contingency operations.*

Combatant commands integrate Combat Mission Forces and Cyber Protection Teams into plans and operations and employ them in cyberspace, while the National Mission Force operates under the Commander of USCYBERCOM. Outside of this construct, teams can also be used to support other missions as required by the Department."

Put simply, CYBERCOM has three core missions: defend DODIN (that's the job of the Cyber Protection Forces, operating under the purview of JFHQ-DoDIN of course); provide operational capabilities to military forces engaged in combat (that's the job of the Combat Mission Forces, who will be operating in practice under the purview of regional combatant commanders); and in special circumstances defend the nation more generally (that's the job of the National Mission Forces).

- Can you identify circumstances that you feel would justify deployment of the National Mission Forces?
- Is that question harder or easier than asking the same thing about a physical-domain threat to the United States, such as a ballistic missile or a military aircraft entering U.S. airspace?

Now, a quick primer on what a "combatant command" is and how CYBERCOM fits into that picture: The traditional organizational structure of the Armed Forces of the United States involved a sharp division into a series of separate "service branches": the Army, Navy, Air Force, and Marines (and the Coast Guard as well, though its precise status is complicated). The several branches not only recruited, trained, and equipped their own forces, but in the past they also planned and commanded their own operations (often, though not always, in coordination with one another). Today, they continue to recruit, train, and equip separately, but they no longer plan and command operations independently. We now have a "joint forces" model for purposes of actual operations. Under this model, assets generated by each branch come under the operational control of a single, unified command structure. More specifically, we now have a globe-spanning series of geographically-defined "combatant commands," such as Central Command (CENTCOM, which encompasses the Middle East through to Afghanistan) and Indo-Pacific Command (INDOPACOM).

So far so good, but it gets more complicated. In addition to these geographically defined commands, we also have several additional commands that have no geographic boundaries but instead are defined by the particular functions they perform or support. CYBERCOM is such a command. Special Operations Command (SOCOM) is another. These functional commands are like the geographic ones in that their subordinate units and personnel are mostly generated by the various service branches, but then brought together under a "joint" command structure for operational purposes. In CYBERCOM's case, that means that Army, Navy, Air Force, and Marine units and personnel make up the various Cyber Mission Forces.

- You likely have heard of the creation of Space Force as a new service branch in the U.S. military. Should we create "Cyber Force," such that the task of recruiting, training, and equipping cyber mission teams falls to an independent branch rather than Army, Navy, Air Force, and Marines?

C. CYBERCOM's role during armed conflict: Joint Task Force-Ares as a case study

Let's start with a word of caution about terminology. People often use phrases like cyber war and cyber attack without meaning to claim that an actual circumstance of war exists or that a particular action amounts to an attack in a war-related sense. Such terms are just convenient shorthands in such cases. Unfortunately, such language can sometimes lead to confusion. Always be clear about your own use of these terms, and be alert to how others might be using (or misusing) them.

- Do a search to find recent news articles that use the words "cyberwar" or "cyber attack." Do any of the examples seem misplaced?

Having said that, let us assume for the moment that we do have a circumstance of armed conflict, such as the conflict in Iraq and Syria involving the Islamic State. What are the critical policy and legal issues that arise when CYBERCOM conducts operations for effect in that setting?

To answer that, let's use the experience of Joint Task Force-Ares as a case study. JTF-Ares is the name of the CYBERCOM mission to engage the Islamic State in the cyber domain, for purposes of both intelligence-gathering and achieving operational effects (either stand-alone operational effects such as shutting down a website operated by IS, or effects intended to support kinetic operations in the field). In 2016, the public began to catch glimpses of internal U.S. government debates associated with these activities, glimpses that highlighted a variety of recurring issues raised by this scenario. The following text is excerpted from a post I wrote for *Lawfare* in 2017, marshaling the public reports that existed up to that point and using them to highlight the issues:

> **1. July 2016 – Reports of DOD frustration over pace of anti-ISIS cyber operations**
>
> In July 2016, the Washington Post (Ellen Nakashima & Missy Ryan) reported on CYBERCOM's efforts to disrupt the Islamic State's online activities (internal

communications, external propaganda, financing, etc.), emphasizing the view of DOD leadership that CYBERCOM was underperforming:

"*An unprecedented Pentagon cyber-offensive against the Islamic State has gotten off to a slow start, officials said, frustrating Pentagon leaders and threatening to undermine efforts to counter the militant group's sophisticated use of technology for recruiting, operations and propaganda. … But defense officials said the command is still working to put the right staff in place and has not yet developed a full suite of malware and other tools tailored to attack an adversary dramatically different from the nation-states Cybercom was created to fight. … Although officials declined to detail current operations, they said that cyberattacks occurring under the new task force might, for instance, disrupt a payment system, identify a communications platform used by Islamic State members and knock it out, or bring down Dabiq, the Islamic State's online magazine.* …"

The report is an excellent snapshot of several distinct challenges the military use of computer network operations can pose.

One such challenge is **operational capacity.** The story suggests that CYBERCOM simply did not have the right personnel and the right exploits on hand for this particular mission, at least at the start. That's a problem that can be fixed, and the report details the steps DOD began taking in 2016 to do just that.

Another challenge is the need to have an effective process for **deconfliction between intelligence-collection and operational-effect equities.** As the article summarized the issue:

"*Whenever the military undertakes a cyber-operation to disrupt a network, the intelligence community may risk losing an opportunity to monitor communications on that network. So military cybersecurity officials have worked to better coordinate their target selection and operations with intelligence officials.*"

This is not a novel tension, in the abstract. For as long as there has been signals intelligence, there have been tensions of this kind. When one side has access to the other's communications, there will always be tension between the temptation to exploit that access for operational effect (with the opportunity cost of risking loss of that access going forward as the enemy realizes it has been monitored) and the temptation to instead exploit it for indirect intelligence advantage (with the opportunity cost of forgoing direct operational advantage in at least some cases). World War II provides famous examples. And so one might fairly ask: is there anything really different about computer network operations, warranting special attention to the topic in this setting?

Perhaps. In this domain there is much more overlap between the means of collection and the means of carrying out a disruptive operations. Indeed, those means often will be the exact same: a particular exploit providing access to an enemy device, network, etc. It seems to me that this ensures that the tension between collection and operational equities will arise with greater frequency, and less room for workarounds, than in more familiar settings.

Having mentioned both the operational capacity concern and the competing-equities concern, now is a good time to emphasize the significance of the status-quo for NSA and CYBERCOM: the dual-hatted commander. Whereas more familiar, traditional scenarios involving tension between collection and operational equities usually involve distinct underlying institutions and commanders, the status quo with respect to computer network

operations has always (well, the past seven years) involved the dual-hatting of NSA's director and CYBERCOM's commander.

This model in theory ensures that neither institution has a home-field advantage, and maximizes the chance that the key decisionmaker (yes, there can be important decisions both below and above the dual-hat, but the dual-hat is obviously in the key position) fully buys into and fully grasps the importance of each institution's mission.

Of course, it is possible that the dual-hat might tilt one direction to an unfair or undesirable degree. And it is possible that some might perceive such a tilt even when there isn't one. As 2016 wore on, questions of this kind began to appear in public, and by September the media was reporting that DNI Clapper and SecDef Carter both were in favor of splitting up the dual-hat. It was not the first time this topic had come up, to be sure; President Obama had considered ordering a split in 2013 (during the aftermath of the Snowden controversy), but had not taken that step at least in part out of concern about CYBERCOM's independent operational capacity. Now the idea appeared to have momentum.

A report from Ellen Nakashima in the Washington Post that same month suggested that this momentum was in part a product of CYBERCOM's operational maturation, but also in significant part driven by the perception that Admiral Rogers, the current dual-hat, favored collection equities to an undue extent:

"*Whether or not it's true, the perception with Secretary Carter and [top aides] has become that the intelligence agency has been winning out at the expense of [cyber] war efforts," said one senior military official….*

(See also this report by the New York Times, stating that frustration along these same lines contributed to the effort to get President Obama to remove Admiral Rogers in late 2016.)

The Washington Post report also highlighted concerns that splitting NSA and CYBERCOM at the leadership level might actually weaken rather than empower CYBERCOM, as NSA inevitably would become free to withhold from CYBERCOM at least some exploits or other forms of access so that sources would not be lost:

"*Cyber Command's mission, their primary focus, is to degrade or destroy," the former official said. "NSA's is exploit [to gather intelligence] only. So without having one person as the leader for both, the bureaucratic walls will go up and you'll find NSA not cooperating with Cyber Command to give them the information they'll need to be successful."*

2. December 2016 – Congress puts on the brakes

Against this backdrop, Congress intervened in late 2016 to slow down the Obama administration's move to split the dual-hat. Section 1642 of the NDAA FY'17, enacted in late December, provides that NSA and CYBERCOM must continue to share a dual-hatted director/commander unless and until the Secretary of Defense and the Chairman of the Joint Chiefs of Staff jointly certify to certain Congressional committees (SASC & HASC; SSCI & HPSCI; and the Appropriations Committees) that separation will not pose "unacceptable" risks to CYBERCOM's effectiveness, and that the following six conditions are met:

"*(i) Robust operational infrastructure has been deployed that is sufficient to meet the unique cyber mission needs of the United States Cyber Command and the National Security Agency, respectively.*

(ii) Robust command and control systems and processes have been established *for planning,* ***deconflicting****, and executing* ***military cyber operations****.*

(iii) The tools and weapons used in cyber operations are sufficient for achieving required effects.

(iv) Capabilities have been established to enable intelligence collection and operational preparation of the environment for cyber operations.

(v) Capabilities have been established to train cyber operations personnel, test cyber capabilities, and rehearse cyber missions.

(vi) The cyber mission force has achieved ***full operational capability****.* Section 1642(b)(2)(C) (emphasis added).

President Obama's signing statement criticized Congress for imposing this requirement, but did not include a claim that it was unconstitutional. It remains the law at this time.

3. Early 2017 – Complications in the War Against the Islamic State

While lawmakers and policymakers wrestled with the pros and cons of splitting NSA and CYBERCOM, computer network operations against the Islamic State continued to accelerate.

Along the way, however, new problems emerged.

As Ellen Nakashima of the Washington Post reported in May 2017, CYBERCOM by late 2016 had encountered a new set of challenges in its enhanced effort to shut down ISIS sites and platforms: **third-country effects.**

"*A secret global operation by the Pentagon late last year to sabotage the Islamic State's online videos and propaganda sparked fierce debate inside the government over whether it was necessary to notify countries that are home to computer hosting services used by the extremist group, including U.S. allies in Europe. … Cybercom developed the campaign under pressure from then-Defense Secretary Ashton B. Carter, who wanted the command to raise its game against the Islamic State. But when the CIA, State Department and FBI got wind of the plan to conduct operations inside the borders of other countries without telling them, officials at the agencies immediately became concerned that the campaign could undermine cooperation with those countries on law enforcement, intelligence and counterterrorism. The issue took the Obama National Security Council weeks to address…*"

This article highlights a third significant challenge associated with computer network operations: attacking the enemy's online presence often requires, or at least risks, some degree of impact on servers located in other countries. Third-country impact involves both legal and policy challenges, and as the quote above illustrates it also brings into play otherwise-unrelated equities of other agencies. Thus, the competing-equities tension is not just a clash between collection and operational equities, but in some cases many others as well. The dual-hat command structure is primarily an answer only to the former, not the latter.

Meanwhile, a sobering reality about the utility of cyberattacks on Islamic State communications began to become clear: the effects often did not last. This was the thrust

> of an important piece by David Sanger and Eric Schmitt in the New York Times in June 2017:
>
> *"[S]ince they began training their arsenal of cyberweapons on …internet use by the Islamic State, the results have been a consistent disappointment, American officials say. … [It] has become clear that recruitment efforts and communications hubs reappear almost as quickly as they are torn down. … "In general, there was some sense of disappointment in the overall ability for cyberoperations to land a major blow against ISIS," or the Islamic State, said Joshua Geltzer, who was the senior director for counterterrorism at the National Security Council until March. "This is just much harder in practice than people think..."*
>
> This suggested that the military equities that some felt had been undervalued by Admiral Rogers in the past were less weighty than proponents had assumed. Nonetheless, momentum towards separation—and concern that the dual-hat unduly favors collection equities—continues.
>
> In mid-July, reports emerged that the Pentagon had submitted to the Trump administration a plan for effectuating the split, with some of the accompanying commentary continuing to advance the argument that NSA holds CYBERCOM back to an improper extent:
>
> *"The goal, [unnamed U.S. officials] said, is to give U.S. Cyber Command more autonomy, freeing it from any constraints that stem from working alongside the NSA, which is responsible for monitoring and collecting telephone, internet and other intelligence data from around the world — a responsibility that can sometimes clash with military operations against enemy forces."*

This account raises a number of questions for you to consider:

- Can you list the variables that may have constrained CYBERCOM in conducting operations for effect against the Islamic State?
- What are the pros and cons of ending the dual-hat arrangement?
- Military operations that produce damage in the physical world often are followed by enemy efforts to repair that damage and restore functionality. Is there reason to think such remediation efforts are, on the whole, easier in cyberspace?
- Can you explain the diplomatic challenge associated with the risk of third-party effects?
- What steps, if any, would you recommend embracing in order to address that challenge?

The account above refers to interagency battles over potential CYBERCOM operations, with CIA, State, and Justice objecting at certain points. This might cause you to wonder: What put those organizations in a position even to know about those plans, let alone to object effectively at the White House level? The answer has to do with an Obama administration policy directive that reportedly required interagency vetting of this sort for military cyber operations expected to have effects outside of areas of active hostilities. Notably, Trump administration officials have announced that this requirement has been revoked (in the form of National Security Presidential Memorandum 13, the precise details of which are not yet public). Some have argued, since then,

that there remains a strong element of interagency vetting in these cases. But let's assume that there is at least a reduction in the extent of that vetting:

- What are the pros and cons of weakening the interagency vetting requirement?

The third-party effect issue also raises a profoundly-sensitive set of international law issues. How so? Here's a thumbnail sketch.

First, Article 2(4) of the U.N. Charter prohibits the "use of force" in international affairs, absent special conditions including (i) a specific type of approval by the U.N. Security Council or (ii) the existence of an armed attack to which this action is a necessary and proportional response. If the U.S. government were to take action in the cyber domain to shut down an Islamic State propaganda site run from a server in Germany, without seeking consent from the German government:

- Would that constitute a "use of force" implicating Article 2(4)?
- Would the U.S. government likely be able to get a U.N. Security Council resolution authorizing the action?
- If this could be done, would timing nonetheless be an issue?
- The United States might argue that it and its coalition partners are responding to an armed attack from the Islamic State. Assume that this is a plausible argument to explain the use of force against the Islamic State in Iraq and Syria. Is it obvious that the same would be true as to an activity taking effect in Germany?

If a cyber-domain activity does not constitute a "use of force" regulated by the U.N. Charter, it might nonetheless constitute an "internationally-wrongful act" in the sense of being a prohibited "intervention" in the internal affairs of a sovereign state. This is a hot topic in international law, currently. To understand why, consider this excerpt from a much-noted articulation of the British view of these questions, from then Attorney General Jeremy Wright in May 2018:

> "In certain circumstances, cyber operations which do not meet the threshold of the use of force but are undertaken by one state against the territory of another state without that state's consent will be considered a breach of international law.
>
> The international law prohibition on intervention in the internal affairs of other states is of particular importance in modern times when technology has an increasing role to play in every facet of our lives, including political campaigns and the conduct of elections. As set out by the International Court of Justice in its judgment in the Nicaragua case, the purpose of this principle is to ensure that all states remain free from external, coercive intervention in the matters of government which are at the heart of a state's sovereignty, such as the freedom to choose its own political, social, economic and cultural system.
> The precise boundaries of this principle are the subject of ongoing debate between states, and not just in the context of cyber space. But the practical application of the principle in this context would be the use by a hostile state of cyber operations to manipulate the electoral system to alter the results of an election in another state, intervention in the

fundamental operation of Parliament, or in the stability of our financial system. Such acts must surely be a breach of the prohibition on intervention in the domestic affairs of states.

Furthermore, a breach of this principle of non-intervention provides victim states with the ability to take action in response that would otherwise be considered unlawful, but which is permissible if it is aimed at returning relations between the hostile state and the victim state to one of lawfulness, and bringing an end to the prior unlawful act. Such action is permissible under the international law doctrine of countermeasures. Put simply, if a hostile state breaches international law as a result of its coercive actions against the target state's sovereign freedoms, then the victim state can take action to compel that hostile state to stop.

Consistent with the de-escalatory nature of international law, there are clear restrictions on the actions that a victim state can take under the doctrine of countermeasures. A countermeasure can only be taken in response to a prior internationally wrongful act committed by a state, and must only be directed towards that state. This means that the victim state must be confident in its attribution of that act to a hostile state before it takes action in response. In cyberspace of course, attribution presents particular challenges, to which I will come in a few moments. Countermeasures cannot involve the use of force, and they must be both necessary and proportionate to the purpose of inducing the hostile state to comply with its obligations under international law.

These restrictions under the doctrine of countermeasures are generally accepted across the international law community. The one area where the UK departs from the excellent work of the International Law Commission on this issue is where the UK is responding to covert cyber intrusion with countermeasures.

In such circumstances, we would not agree that we are always legally obliged to give prior notification to the hostile state before taking countermeasures against it. The covertness and secrecy of the countermeasures must of course be considered necessary and proportionate to the original illegality, but we say it could not be right for international law to require a countermeasure to expose highly sensitive capabilities in defending the country in the cyber arena, as in any other arena.

In addition, it is also worth stating that, as a matter of law, there is no requirement in the doctrine of countermeasures for a response to be symmetrical to the underlying unlawful act. What matters is necessity and proportionality, which means that the UK could respond to a cyber intrusion through non-cyber means, and vice versa.

Through the principle of non-intervention, it is clear that the international community has set a boundary at which interference in another state's sovereign freedoms is considered internationally wrongful and as such, in breach of international law, giving rise to the right to take action which may otherwise be unlawful in response. As I have already mentioned, the precise parameters of this principle remain the subject of ongoing debate in the international law community, but a further contested area amongst those engaged in the application of international law to cyber space is the regulation of activities that fall below the threshold of a prohibited intervention, but nonetheless may be perceived as affecting the territorial sovereignty of another state without that state's prior consent.

Some have sought to argue for the existence of a cyber specific rule of a "violation of territorial sovereignty" in relation to interference in the computer networks of another state without its consent.

> Sovereignty is of course fundamental to the international rules-based system. But I am not persuaded that we can currently extrapolate from that general principle a specific rule or additional prohibition for cyber activity beyond that of a prohibited intervention. The UK Government's position is therefore that there is no such rule as a matter of current international law."

It is widely thought that this understanding of the legal status of sovereignty reflects the U.S. position as well. Here is language from a speech by Brian Egan, then the Legal Adviser of the State Department, in 2016:

> "In certain circumstances, one State's non-consensual cyber operation in another State's territory could violate international law, even if it falls below the threshold of a use of force. This is a challenging area of the law that raises difficult questions. The very design of the Internet may lead to some encroachment on other sovereign jurisdictions. Precisely when a non-consensual cyber operation violates the sovereignty of another State is a question lawyers within the U.S. government continue to study carefully, and it is one that ultimately will be resolved through the practice and opinio juris of States.
>
> Relatedly, consider the challenges we face in clarifying the international law prohibition on unlawful intervention. As articulated by the International Court of Justice (ICJ) in its judgment on the merits in the Nicaragua Case, this rule of customary international law forbids States from engaging in coercive action that bears on a matter that each State is entitled, by the principle of State sovereignty, to decide freely, such as the choice of a political, economic, social, and cultural system. This is generally viewed as a relatively narrow rule of customary international law, but States' cyber activities could run afoul of this prohibition. For example, a cyber operation by a State that interferes with another country's ability to hold an election or that manipulates another country's election results would be a clear violation of the rule of non-intervention. For increased transparency, States need to do more work to clarify how the international law on non-intervention applies to States' activities in cyberspace.

- Can you explain how the varying national interests and circumstances of the United States, Russia, and China might cause them to take different positions on these issues?

25. Government Hacking: Grey-Zone Competition

Cyber domain operations *during* armed conflict raise interesting legal and policy questions, plainly. Yet they are not quite as interesting, perhaps, as the ones raised by interstate competition carried out in the cyber domain *below* the threshold of armed conflict.

- Can you think of some prominent recent examples that we have discussed, involving actions by other states that were conducted in the cyber domain against U.S. interests?
- Can you think of any examples of the United States conducting such operations against other states?
- *Should* the United States ever conduct such operations (excluding espionage)?
- If the United States decided not to conduct such operations for effect outside the context of armed conflict, does it follow that the United States would not attempt to hack foreign systems before a conflict occurred?

Apart from the policy questions associated with competition below the threshold of armed conflict, there are also complex legal questions.

First, and most obviously, all of the same questions raised above in the armed-conflict context regarding the U.N. Charter, the prohibition on intervention, and the idea of a general rule of sovereignty also come into play in this non-conflict setting.

- Do you think that Russia's combined hacking-and-disinformation operation against the United States in 2016 violated international law?
- Why, if at all, would it matter if the answer to that question is yes?

Second, this scenario raises a tricky question of U.S. domestic law when CYBERCOM is asked to perform the operation in question. Why? Because until recently there was uncertainty regarding whether CYBERCOM has affirmative authority under U.S. law to take on this particular role (an issue that would not arise if we were speaking of a circumstance of armed conflict).

In recent years, this question had become a source of considerable friction within the government, for better or worse. But Congress then took action (in the most-recent National Defense Authorization Act) to prune away this concern.

Section 1642 of the NDAA Fiscal Year '19 provides in relevant part that:

> **(a) AUTHORITY TO DISRUPT, DEFEAT, AND DETER CYBER ATTACKS.—**
>
> **(1) IN GENERAL.—**In the event that the National Command Authority determines that the Russian Federation, People's Republic of China, Democratic People's Republic of Korea, or Islamic Republic of Iran is conducting an active, systematic, and ongoing campaign of attacks against the Government or people of the United States in cyberspace, including attempting to influence American elections and democratic political processes, the National Command Authority may authorize the Secretary of Defense, acting through the Commander of the United States Cyber Command, to take appropriate and proportional action in foreign cyberspace to disrupt, defeat, and deter such attacks under the authority and policy of the Secretary of Defense to conduct cyber operations and information operations as traditional military activities.

- Can you break this statute down into "elements" that must be satisfied?
- What is the consequence if one or more of these elements is not satisfied?
- Can you relate the likely impact of this authority to the likely impact of the decision by the Trump administration (discussed earlier) to reduce interagency vetting of military cyber operations expected to have effects outside of combat zones?

Going deeper…

Congratulations, you've reached the end of the book! I'm grateful to you for reading it, and I hope you feel your time and energy were well spent. Please do not hesitate to email me (rchesney@law.utexas.edu) with any thoughts or suggestions you care to share; I'll be happy to hear from you. Meanwhile, if you want to stay engaged with this topic, I have some recommendations for you.

First, two cybersecurity-focused podcasts:

The Steptoe Cyberlaw Podcast - a weekly review of key *legal* developments relating to cybersecurity, followed by in-depth interviews with a wide variety of guests.

Risky Business – a weekly review of important news from the *technology* side of cybersecurity, also followed by in-depth interviews with a variety of guests

Second, two projects of mine that sometimes touch on cybersecurity issues:

The National Security Law Podcast – Each week, my UT-Austin colleague Steve Vladeck and I review, discuss, and debate the latest national security law developments (always with a healthy dose of frivolity; the show is *not* for those who prefer their national security law discussions to be dry and serious at all times).

Lawfare – Lawfare is the nation's leading source of online analysis of current national security legal issues. Co-founded by Ben Wittes, Jack Goldsmith, and me, Lawfare covers an array of topics including those associated with cyber domain activities.

You also should follow the daily developments in this space using resources such as Politico's Morning Cybersecurity and the Washington Post feature The Cybersecurity 202.

The end.