Secure Software Engineering and Management

UPenn CIS 7000-010

Michael Hicks

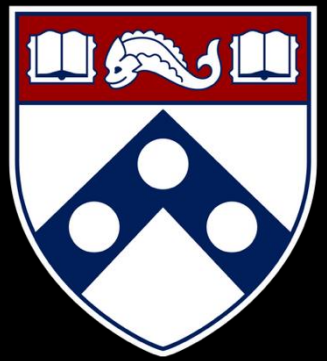

Cybersecurity Regulation and Compliance

Misaligned Incentives

Recall from Anderson:

- Markets systematically **under-supply** security
- **Incentives are misaligned** — externalities, information asymmetry
- The party that pays for security isn't always the one that suffers its absence

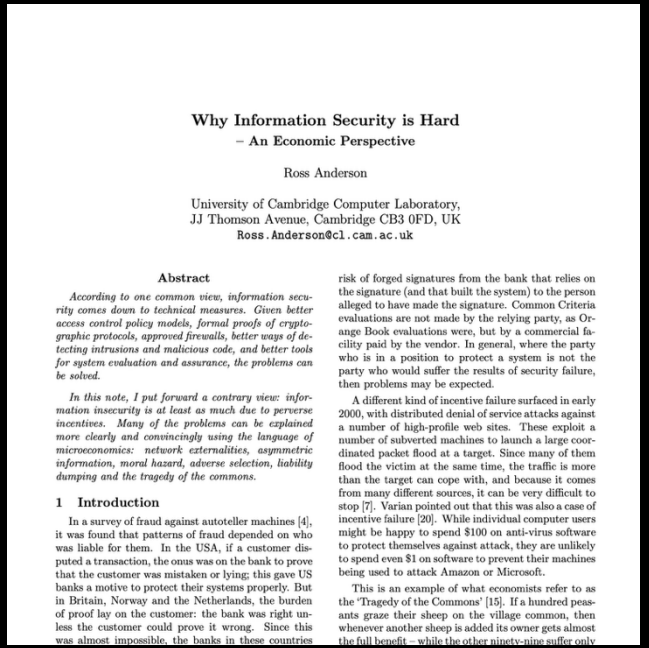

Why Information Security is Hard
– An Economic Perspective

Ross Anderson

University of Cambridge Computer Laboratory,
JJ Thomson Avenue, Cambridge CB3 0FD, UK
Ross.Anderson@cl.cam.ac.uk

Abstract

According to one common view, information security comes down to technical measures. Given better access control policy models, formal proofs of cryptographic protocols, approved firewalls, better ways of detecting intrusions and malicious code, and better tools for system evaluation and assurance, the problems can be solved.

In this note, I put forward a contrary view: information insecurity is at least as much due to perverse incentives. Many of the problems can be explained more clearly and convincingly using the language of microeconomics: network externalities, asymmetric information, moral hazard, adverse selection, liability dumping and the tragedy of the commons.

1 Introduction

In a survey of fraud against autoteller machines [4], it was found that patterns of fraud depended on who was liable for them. In the USA, if a customer disputed a transaction, the onus was on the bank to prove that the customer was mistaken or lying; this gave US banks a motive to protect their systems properly. But in Britain, Norway and the Netherlands, the burden of proof lay on the customer: the bank was right unless the customer could prove it wrong. Since this was almost impossible, the banks in these countries risk of forged signatures from the bank that relies on the signature (and that built the system) to the person alleged to have made the signature. Common Criteria evaluations are not made by the relying party, as Orange Book evaluations were, but by a commercial facility paid by the vendor. In general, where the party who is in a position to protect a system is not the party who would suffer the results of security failure, then problems may be expected.

A different kind of incentive failure surfaced in early 2000, with distributed denial of service attacks against a number of high-profile web sites. These exploit a number of subverted machines to launch a large coordinated packet flood at a target. Since many of them flood the victim at the same time, the traffic is more than the target can cope with, and because it comes from many different sources, it can be very difficult to stop [7]. Varian pointed out that this was also a case of incentive failure [20]. While individual computer users might be happy to spend $100 on anti-virus software to protect themselves against attack, they are unlikely to spend even $1 on software to prevent their machines being used to attack Amazon or Microsoft.

This is an example of what economists refer to as the 'Tragedy of the Commons' [15]. If a hundred peasants graze their sheep on the village common, then whenever another sheep is added its owner gets almost the full benefit – while the other ninety-nine suffer only

Regulation is the classical response to market failure.

Today: does cyber regulation actually align the incentives?

The Big Question Today

If I am building a cloud-hosted service with valuable customer data:

- **What does the law require me to do?**
- **Does doing it make me more secure?**
- And if not — **why not, and what would work better?**

Outline

1. Case study: Equifax and FTC enforcement
2. Compliance ≠ security: the debate
3. The US regulatory toolkit
4. The EU approach: horizontal regulation
5. Incident reporting and the Cyber NTSB

Part 1: Case Study

Equifax Says Cyberattack May Have Affected 143 Million in the U.S.

By Tara Siegel Bernard, Tiffany Hsu, Nicole Perlroth and Ron Lieber
Sept. 7, 2017

Equifax, one of the three major consumer credit reporting agencies, said on Thursday that hackers had gained access to company data that potentially compromised sensitive information for 143 million American consumers, including Social Security numbers and driver's license numbers.

The attack on the company represents one of the largest risks to personally sensitive information in recent years, and is the third major cybersecurity threat for the agency since 2015.

Equifax, based in Atlanta, is a particularly tempting target for hackers. If identity thieves wanted to hit one place to grab all the data needed to do the most damage, they would go straight to one of the three major credit reporting agencies.

"This is about as bad as it gets," said Pamela Dixon, executive director of the World Privacy Forum, a nonprofit research group. "If you have a credit report, chances are you may be in this breach. The chances are much better than 50 percent."

Equifax 2017: A “Compliant” Company

Credit bureau holding data on ~200M US consumers.

Credentials at the time of breach:

- **PCI DSS Level 1** — annual on-site QSA audit
- **SOC 1 Type II** — financial controls audited annually
- Internal security org + CISO + $250M+ annual security budget

By every standard of the time: a compliant organization.

Timeline of the Breach

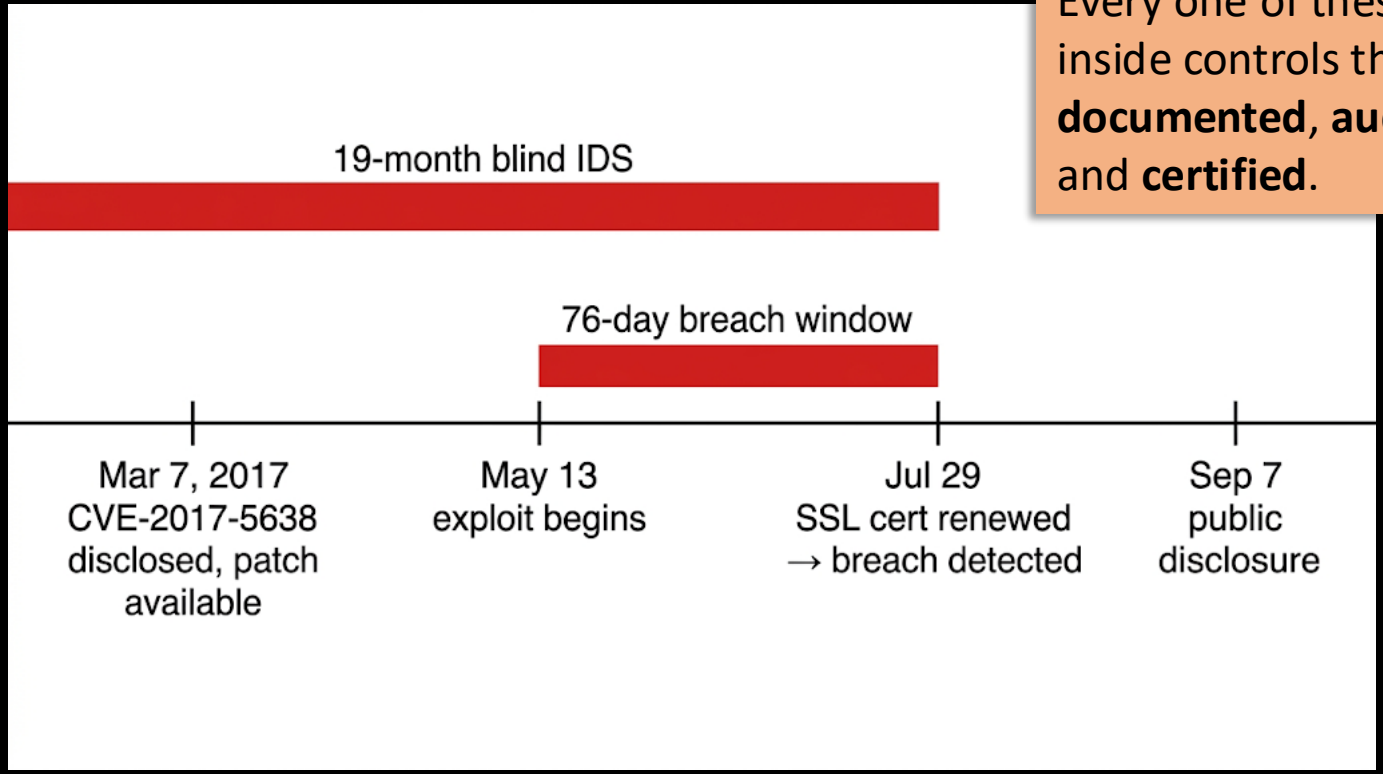

Every one of these sat inside controls that were **documented**, **audited**, and **certified**.

The Breach

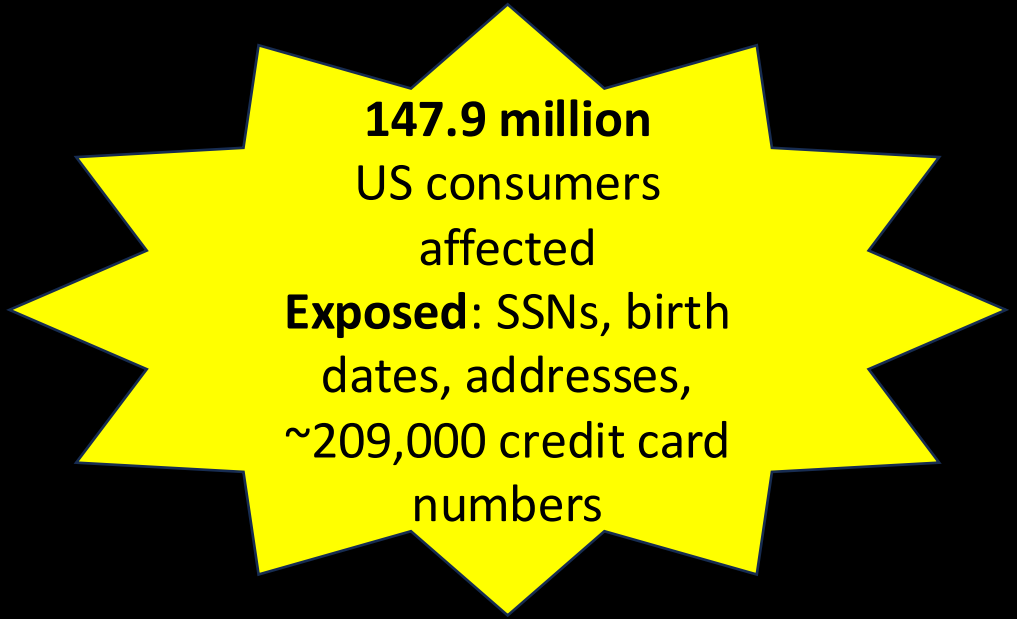
147.9 million
US consumers affected
Exposed: SSNs, birth dates, addresses, ~209,000 credit card numbers

The Breach

40 days

Breach discovered, Jul 29

Public disclosure Sep 7

147.9 million US consumers affected

Exposed: SSNs, birth dates, addresses, ~209,000 credit card numbers

4 days

Incident determined material

Public disclosure

CIO Jun Ying sold ~**$950K** in stock before disclosure → 4 months federal prison

The Fallout: $1.4B by 2019

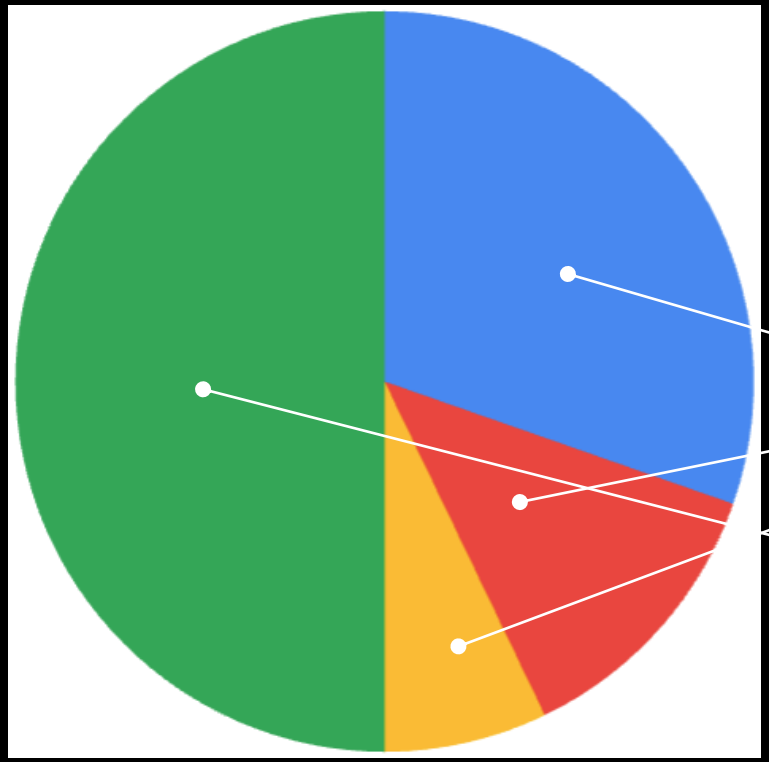

- 50-state AG coordinated action
- Congressional hearings (4 committees, Oct 2017)
- **July 2019**: FTC settlement of **$700M**
  - $425M consumer restitution
  - $175M to states
  - $100M CFPB civil penalty
- Plus **~$700M** incident response, tech, data security upgrades, investigations, lawsuits, …

Remedy Foundation: FTC v. Wyndham (2015)

The question: does the FTC have authority to regulate cybersecurity at all?

Wyndham's argument: no — Congress never passed a specific cybersecurity statute. Section 5 "unfairness" doesn't reach data-security practices. Holding us liable without explicit rules violates due process.

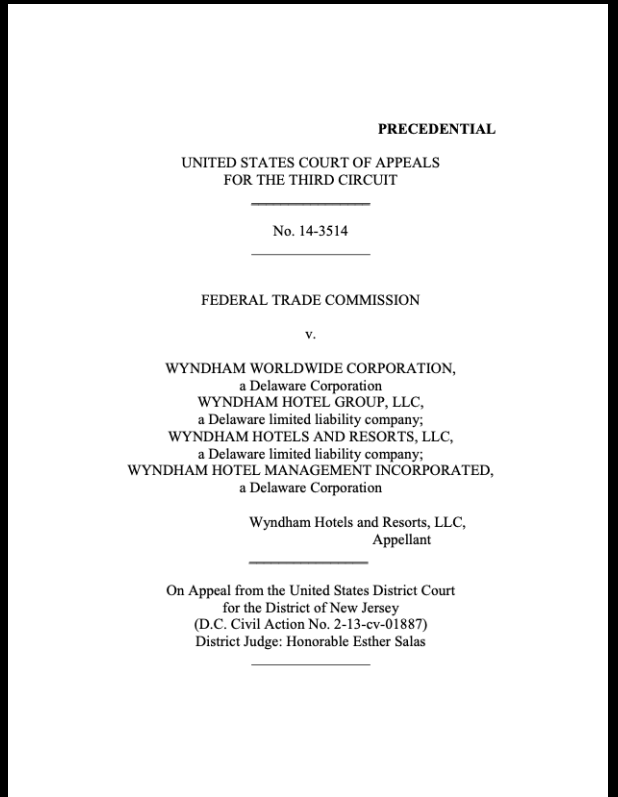

PRECEDENTIAL

UNITED STATES COURT OF APPEALS
FOR THE THIRD CIRCUIT

No. 14-3514

FEDERAL TRADE COMMISSION

v.

WYNDHAM WORLDWIDE CORPORATION,
a Delaware Corporation
WYNDHAM HOTEL GROUP, LLC,
a Delaware limited liability company;
WYNDHAM HOTELS AND RESORTS, LLC,
a Delaware limited liability company;
WYNDHAM HOTEL MANAGEMENT INCORPORATED,
a Delaware Corporation

Wyndham Hotels and Resorts, LLC,
Appellant

On Appeal from the United States District Court
for the District of New Jersey
(D.C. Civil Action No. 2-13-cv-01887)
District Judge: Honorable Esther Salas

The Ruling

Third Circuit, August 2015:

1. Section 5 unfairness **does** encompass data security
2. FTC does **not** need to issue formal cybersecurity regulations first
3. Prior consent decrees + published guidance + complaints gave **fair notice**
4. Wyndham's failures were so basic (cleartext credit cards, no firewalls) that fair notice was easily satisfied

Settlement: 20-year consent order, annual PCI DSS audits, franchisee network safeguards

Why Wyndham Matters

- US (still) has **no** economy-wide cybersecurity statute
- Wyndham gave the FTC de facto cybersecurity authority via Section 5
- Standards built **case by case** through enforcement
- Distinctly American: regulator accretes doctrine rather than Congress legislating

Every subsequent FTC cyber case sits on top of Wyndham

THE FTC'S CASE-BY-CASE DOCTRINE ACCRETION IN DATA SECURITY

A Historical Progression of Legal Precedents

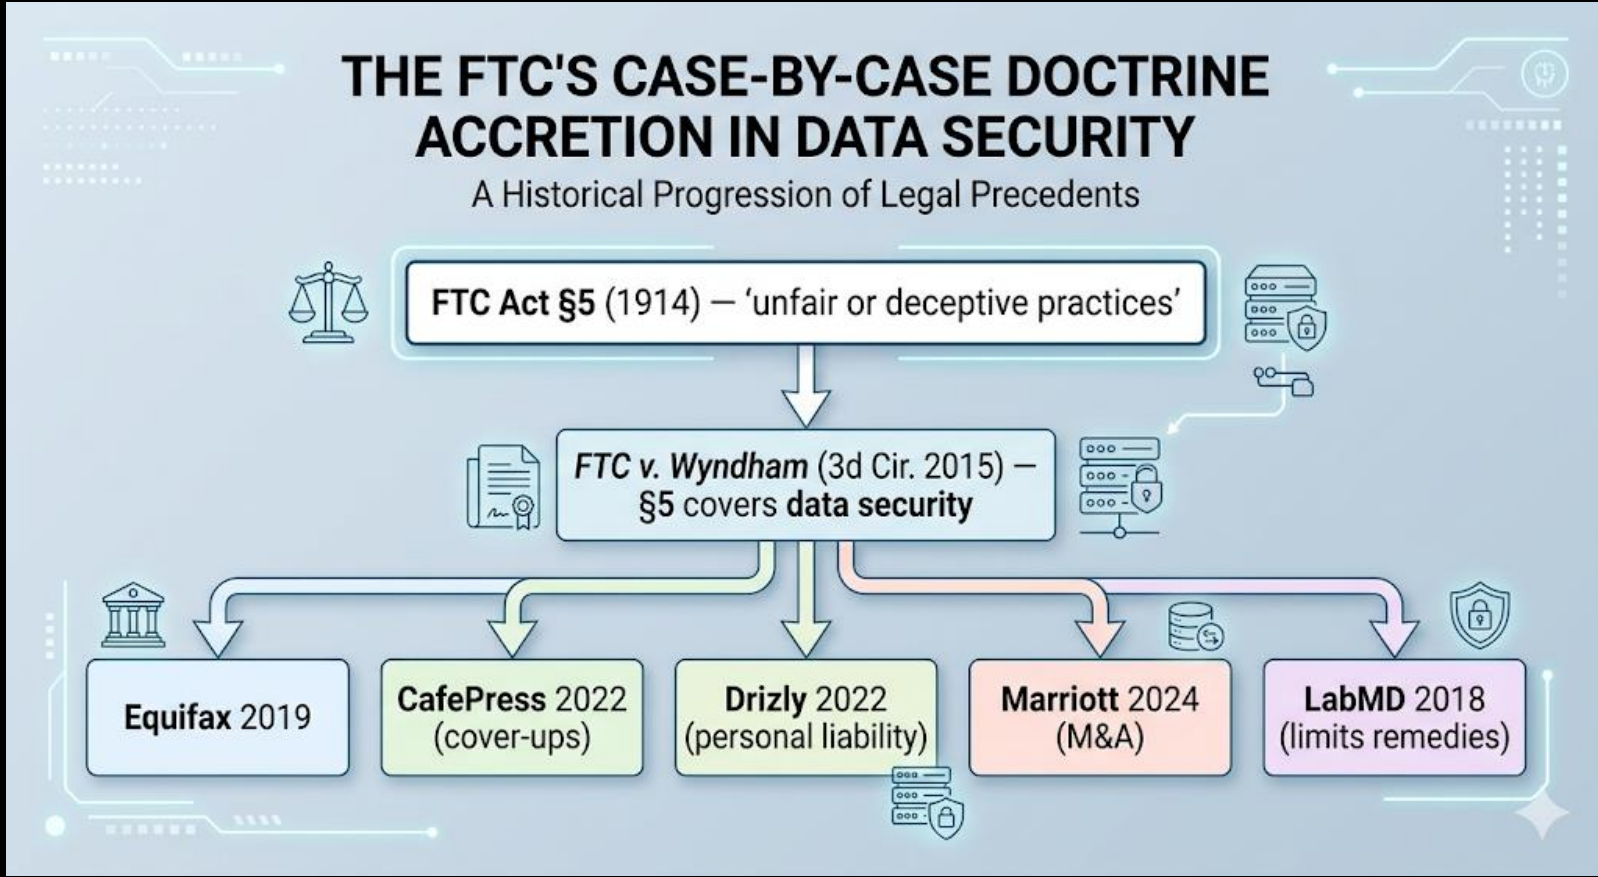

FTC as Cyber Enforcer: A Pattern

Case	Year	Violation	Enforcement
Wyndham	2015	619K cards, 3 breaches	20-yr consent order; PCI audits
Equifax	2019	147M records	$700M settlement (FTC + states + CFPB)
CafePress	2022	22M accounts	$500K penalty; **first cover-up case**
Drizly	2022	2.5M consumers	**CEO bound 10 yrs at any future firm**
Marriott	2024	344M guest records, 3 breaches	20-yr program of indep. security reviews; $52M multistate settlement

State AGs active too: Capital One ($190M, 2021), T-Mobile ($31.5M, 2024), Home Depot ($17.5M, 2020).

Two Pillars of FTC Cyber Authority

Wyndham established the *general* pillar. A *sector-specific* pillar existed.

GLBA = 1999 Gramm-Leach-Bliley Act (aka the Financial Services Modernization Act)

Pillar	Source	Scope	Standard
General	FTC Act §5 + *Wyndham*	All commerce	Case-by-case "unfairness"
Sector-specific	**GLBA Safeguards Rule** (2003)	Non-bank financial institutions	"Reasonable" security program

Wyndham (hotel chain) tested Pillar 1.

Equifax (credit bureau) was caught by *both* — and the $700M settlement invoked Section 5 *and* GLBA Safeguards violations together.

GLBA Safeguards Revision: Rules Strengthened

2021 rule, effective **June 2023**.
Major shift: principles → **prescriptive**.

- Mandatory **encryption** in transit and at rest
- **MFA** for customer data access
- Designated **Qualified Individual** (personal accountability)
- Written incident response plans
- Periodic **pen testing** and vulnerability assessments
- **Annual compliance reporting to boards**

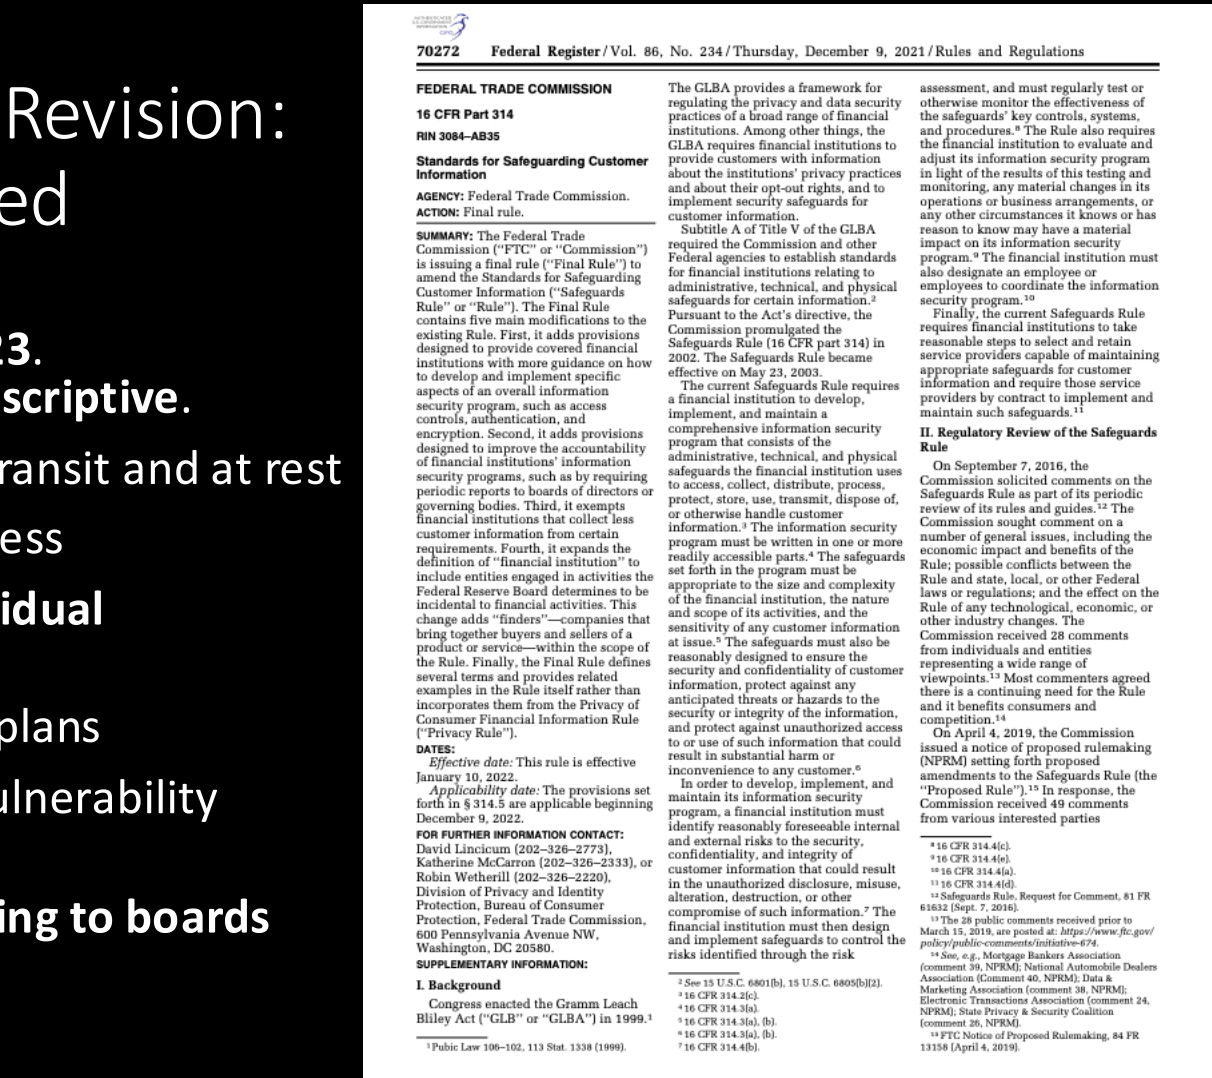

FEDERAL TRADE COMMISSION

16 CFR Part 314

RIN 3084–AB35

Standards for Safeguarding Customer Information

AGENCY: Federal Trade Commission.

ACTION: Final rule.

SUMMARY: The Federal Trade Commission ("FTC" or "Commission") is issuing a final rule ("Final Rule") to amend the Standards for Safeguarding Customer Information ("Safeguards Rule" or "Rule"). The Final Rule contains five main modifications to the existing Rule. First, it adds provisions designed to provide covered financial institutions with more guidance on how to develop and implement specific aspects of an overall information security program, such as access controls, authentication, and encryption. Second, it adds provisions designed to improve the accountability of financial institutions' information security programs, such as by requiring periodic reports to boards of directors or governing bodies. Third, it exempts financial institutions that collect less customer information from certain requirements. Fourth, it expands the definition of "financial institution" to include entities engaged in activities the Federal Reserve Board determines to be incidental to financial activities. This change adds "finders"—companies that bring together buyers and sellers of a product or service—within the scope of the Rule. Finally, the Final Rule defines several terms and provides related examples in the Rule itself rather than incorporates them from the Privacy of Consumer Financial Information Rule ("Privacy Rule").

DATES:

Effective date: This rule is effective January 10, 2022.

Applicability date: The provisions set forth in § 314.5 are applicable beginning December 9, 2022.

FOR FURTHER INFORMATION CONTACT: David Lincicum (202–326–2773), Katherine McCarron (202–326–2333), or Robin Wetherill (202–326–2220), Division of Privacy and Identity Protection, Bureau of Consumer Protection, Federal Trade Commission, 600 Pennsylvania Avenue NW, Washington, DC 20580.

SUPPLEMENTARY INFORMATION:

I. Background

Congress enacted the Gramm Leach Bliley Act ("GLB" or "GLBA") in 1999.[1] The GLBA provides a framework for regulating the privacy and data security practices of a broad range of financial institutions. Among other things, the GLBA requires financial institutions to provide customers with information about the institutions' privacy practices and about their opt-out rights, and to implement security safeguards for customer information.

Subtitle A of Title V of the GLBA required the Commission and other Federal agencies to establish standards for financial institutions relating to administrative, technical, and physical safeguards for certain information.[2] Pursuant to the Act's directive, the Commission promulgated the Safeguards Rule (16 CFR part 314) in 2002. The Safeguards Rule became effective on May 23, 2003.

The current Safeguards Rule requires a financial institution to develop, implement, and maintain a comprehensive information security program that consists of the administrative, technical, and physical safeguards the financial institution uses to access, collect, distribute, process, protect, store, use, transmit, dispose of, or otherwise handle customer information.[3] The information security program must be written in one or more readily accessible parts.[4] The safeguards set forth in the program must be appropriate to the size and complexity of the financial institution, the nature and scope of its activities, and the sensitivity of any customer information at issue.[5] The safeguards must also be reasonably designed to ensure the security and confidentiality of customer information, protect against any anticipated threats or hazards to the security or integrity of the information, and protect against unauthorized access to or use of such information that could result in substantial harm or inconvenience to any customer.[6]

In order to develop, implement, and maintain its information security program, a financial institution must identify reasonably foreseeable internal and external risks to the security, confidentiality, and integrity of customer information that could result in the unauthorized disclosure, misuse, alteration, destruction, or other compromise of such information.[7] The financial institution must then design and implement safeguards to control the risks identified through the risk assessment, and must regularly test or otherwise monitor the effectiveness of the safeguards' key controls, systems, and procedures.[8] The Rule also requires the financial institution to evaluate and adjust its information security program in light of the results of this testing and monitoring, any material changes in its operations or business arrangements, or any other circumstances it knows or has reason to know may have a material impact on its information security program.[9] The financial institution must also designate an employee or employees to coordinate the information security program.[10]

Finally, the current Safeguards Rule requires financial institutions to take reasonable steps to select and retain service providers capable of maintaining appropriate safeguards for customer information and require those service providers by contract to implement and maintain such safeguards.[11]

II. Regulatory Review of the Safeguards Rule

On September 7, 2016, the Commission solicited comments on the Safeguards Rule as part of its periodic review of its rules and guides.[12] The Commission sought comment on a number of general issues, including the economic impact and benefits of the Rule; possible conflicts between the Rule and state, local, or other Federal laws or regulations; and the effect on the Rule of any technological, economic, or other industry changes. The Commission received 28 comments from individuals and entities representing a wide range of viewpoints.[13] Most commenters agreed there is a continuing need for the Rule and it benefits consumers and competition.[14]

On April 4, 2019, the Commission issued a notice of proposed rulemaking (NPRM) setting forth proposed amendments to the Safeguards Rule (the "Proposed Rule").[15] In response, the Commission received 49 comments from various interested parties

[1] Pubic Law 106–102, 113 Stat. 1338 (1999).

[2] *See* 15 U.S.C. 6801(b), 15 U.S.C. 6805(b)(2).

[3] 16 CFR 314.2(c).

[4] 16 CFR 314.3(a).

[5] 16 CFR 314.3(a), (b).

[6] 16 CFR 314.3(a), (b).

[7] 16 CFR 314.4(b).

[8] 16 CFR 314.4(c).

[9] 16 CFR 314.4(e).

[10] 16 CFR 314.4(a).

[11] 16 CFR 314.4(d).

[12] Safeguards Rule, Request for Comment, 81 FR 61632 (Sept. 7, 2016).

[13] The 28 public comments received prior to March 15, 2019, are posted at: *https://www.ftc.gov/policy/public-comments/initiative-674.*

[14] *See, e.g.*, Mortgage Bankers Association (comment 39, NPRM); National Automobile Dealers Association (Comment 40, NPRM); Data & Marketing Association (comment 38, NPRM); Electronic Transactions Association (comment 24, NPRM); State Privacy & Security Coalition (comment 26, NPRM).

[15] FTC Notice of Proposed Rulemaking, 84 FR 13158 (April 4, 2019).

Part 2: Compliance ≠ Security — The Debate

Stepping Back

Equifax

- held multiple certifications
- had a CISO
- spent nine figures on security

And was breached through a **patching failure** and an **expired certificate**.

Post-Equifax, regulators tightened specific rules. But:

Does compliance actually equal security? And if not, why not?

Equifax Is Not an Anomaly

Equifax Is Not an Anomaly

Target (2013):

- Was **PCI-DSS certified** when breached
- Attackers entered via HVAC vendor
- **FireEye alerts ignored**
- 40M cards + 70M customer records
- Cost: $290M+

Heartland Payment Systems (2008), Anthem (2015), Marriott (2018):
all certified compliant at time of breach.

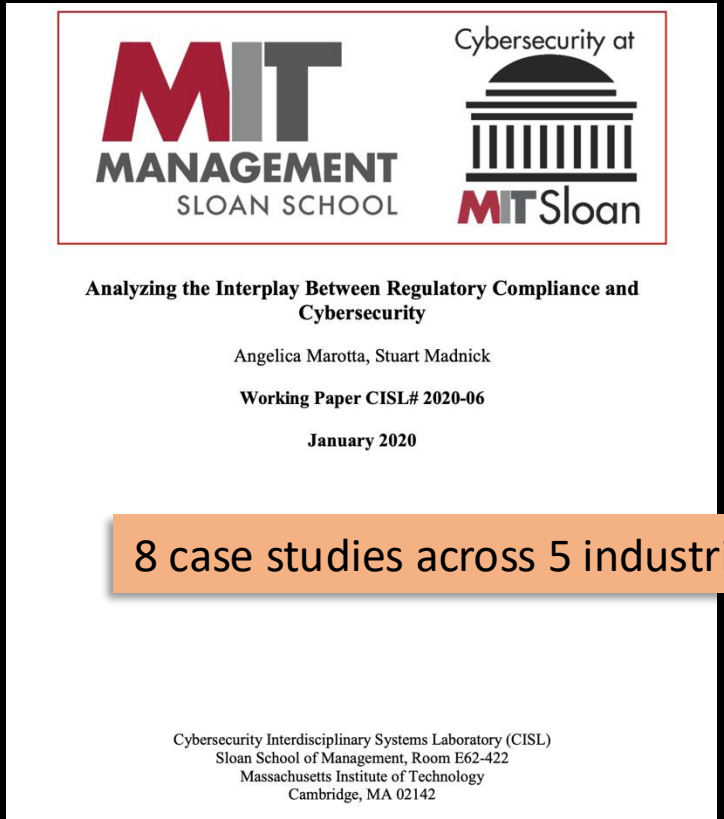
MIT MANAGEMENT SLOAN SCHOOL

Cybersecurity at MIT Sloan

Analyzing the Interplay Between Regulatory Compliance and Cybersecurity

Angelica Marotta, Stuart Madnick

Working Paper CISL# 2020-06

January 2020

8 case studies across 5 industries

Cybersecurity Interdisciplinary Systems Laboratory (CISL)
Sloan School of Management, Room E62-422
Massachusetts Institute of Technology
Cambridge, MA 02142

Compliance Can *Substitute* for Security

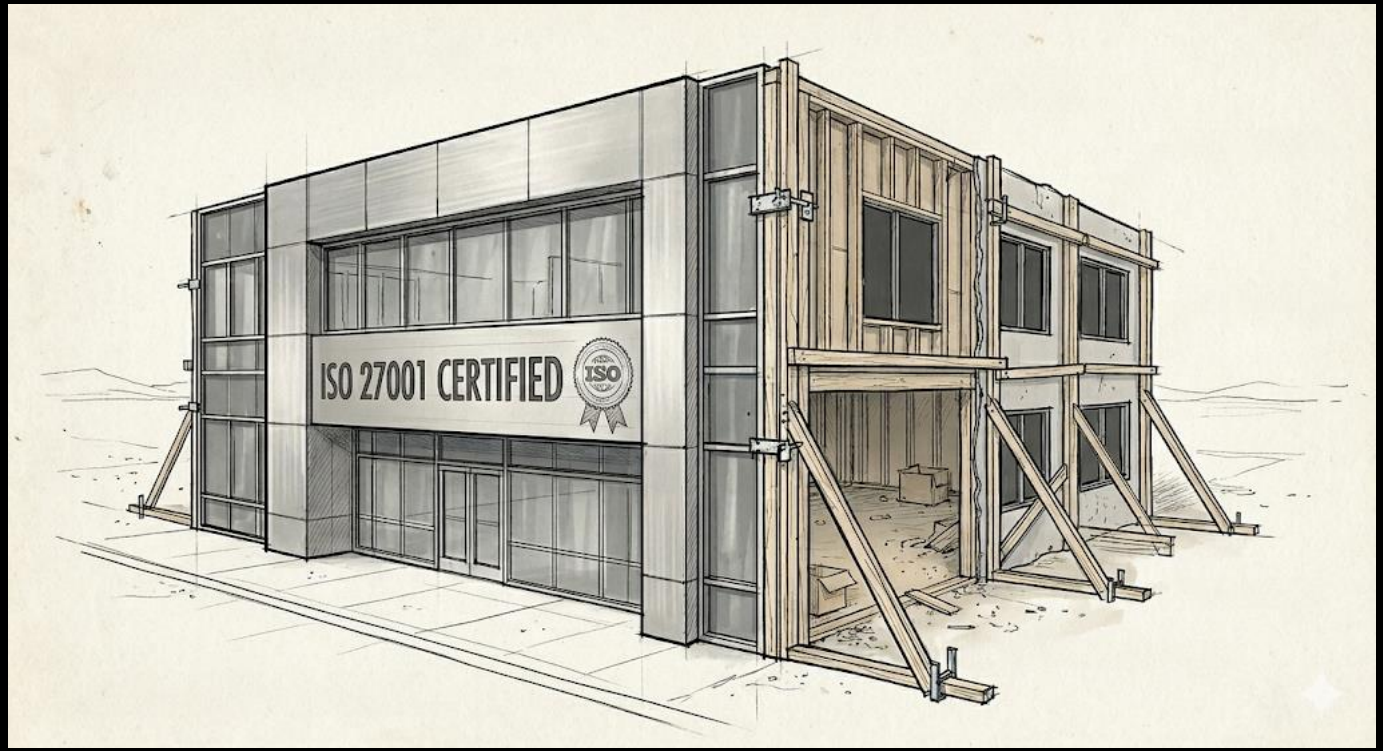

Compliance Can *Substitute* for Security

Optimizing for audit passage ≠ reducing risk.

- Compliance becomes a **ceiling**, not a floor
- “Paper security”: documentation substitutes for operation
- Budget crowded out: compliance work is expensive, visible, measurable
- **Grigg’s silver bullets, redux**: compliance certifications signal “doing security”

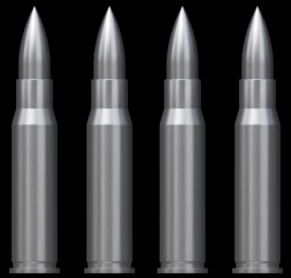

The Compliance Theater

Thesis: cybersecurity compliance has become **bureaucratic theater** that prioritizes documentation and procedural adherence over actual security outcomes.

> "A compliance framework measures the presence of security controls and the quality of documentation. It does not directly measure security."
>
> --Nick Kelly, SecureFlag, 2025

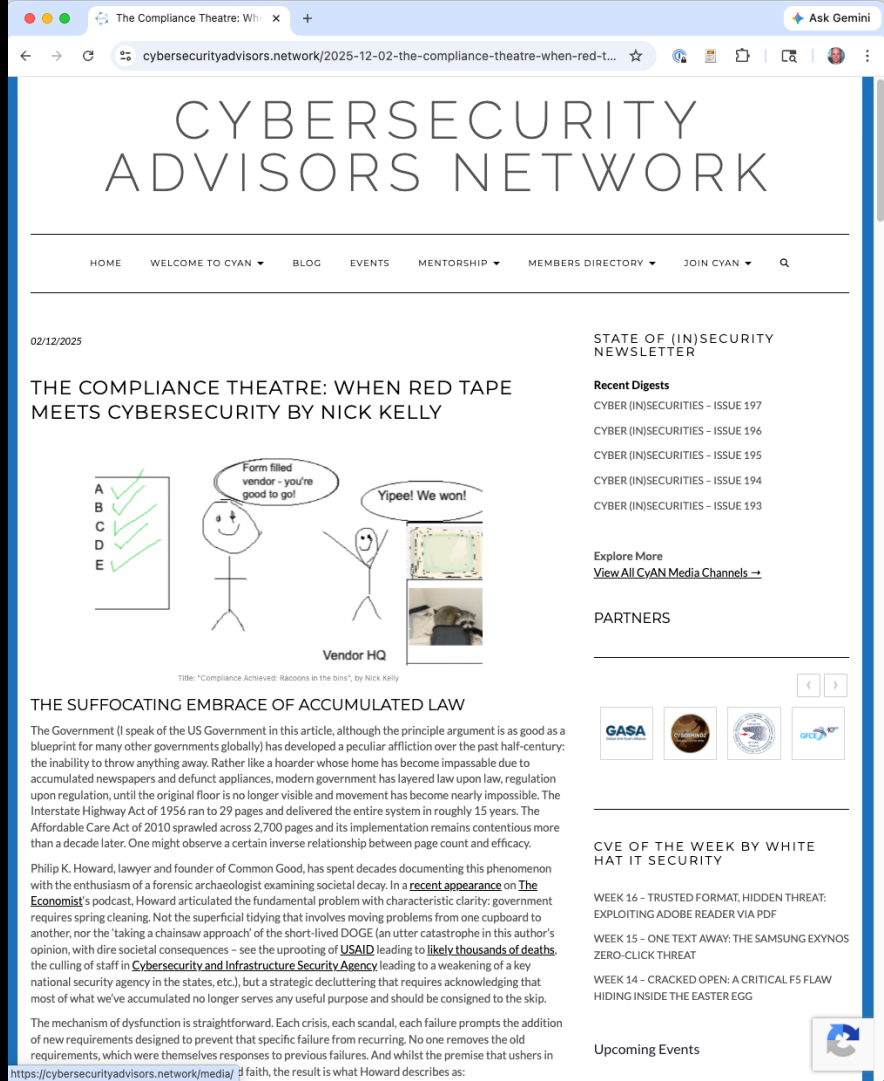

The Compliance Theatre: Wh

cybersecurityadvisors.network/2025-12-02-the-compliance-theatre-when-red-t...

CYBERSECURITY ADVISORS NETWORK

HOME | WELCOME TO CYAN | BLOG | EVENTS | MENTORSHIP | MEMBERS DIRECTORY | JOIN CYAN

02/12/2025

THE COMPLIANCE THEATRE: WHEN RED TAPE MEETS CYBERSECURITY BY NICK KELLY

Title: "Compliance Achieved: Racoons in the bins", by Nick Kelly

THE SUFFOCATING EMBRACE OF ACCUMULATED LAW

The Government (I speak of the US Government in this article, although the principle argument is as good as a blueprint for many other governments globally) has developed a peculiar affliction over the past half-century: the inability to throw anything away. Rather like a hoarder whose home has become impassable due to accumulated newspapers and defunct appliances, modern government has layered law upon law, regulation upon regulation, until the original floor is no longer visible and movement has become nearly impossible. The Interstate Highway Act of 1956 ran to 29 pages and delivered the entire system in roughly 15 years. The Affordable Care Act of 2010 sprawled across 2,700 pages and its implementation remains contentious more than a decade later. One might observe a certain inverse relationship between page count and efficacy.

Philip K. Howard, lawyer and founder of Common Good, has spent decades documenting this phenomenon with the enthusiasm of a forensic archaeologist examining societal decay. In a recent appearance on The Economist's podcast, Howard articulated the fundamental problem with characteristic clarity: government requires spring cleaning. Not the superficial tidying that involves moving problems from one cupboard to another, nor the 'taking a chainsaw approach' of the short-lived DOGE (an utter catastrophe in this author's opinion, with dire societal consequences – see the uprooting of USAID leading to likely thousands of deaths, the culling of staff in Cybersecurity and Infrastructure Security Agency leading to a weakening of a key national security agency in the states, etc.), but a strategic decluttering that requires acknowledging that most of what we've accumulated no longer serves any useful purpose and should be consigned to the skip.

The mechanism of dysfunction is straightforward. Each crisis, each scandal, each failure prompts the addition of new requirements designed to prevent that specific failure from recurring. No one removes the old requirements, which were themselves responses to previous failures. And whilst the premise that ushers in

https://cybersecurityadvisors.network/media/ d faith, the result is what Howard describes as:

STATE OF (IN)SECURITY NEWSLETTER

Recent Digests

CYBER (IN)SECURITIES – ISSUE 197

CYBER (IN)SECURITIES – ISSUE 196

CYBER (IN)SECURITIES – ISSUE 195

CYBER (IN)SECURITIES – ISSUE 194

CYBER (IN)SECURITIES – ISSUE 193

Explore More

View All CyAN Media Channels →

PARTNERS

CVE OF THE WEEK BY WHITE HAT IT SECURITY

WEEK 16 – TRUSTED FORMAT, HIDDEN THREAT: EXPLOITING ADOBE READER VIA PDF

WEEK 15 – ONE TEXT AWAY: THE SAMSUNG EXYNOS ZERO-CLICK THREAT

WEEK 14 – CRACKED OPEN: A CRITICAL F5 FLAW HIDING INSIDE THE EASTER EGG

Upcoming Events

1: The Checkbox Illusion

- Compliance questionnaires measure **whether firms know the right answers**
- They don’t measure **whether implementations work**
- Certified organizations get breached routinely
  - Heartland Payment Systems
  - Anthem
  - Telemessage (2025, federally approved)

Perfect compliance ≠ adequate security.

2: Perverse Incentives

What CISOs are rewarded for:

- ✅ FedRAMP authorization
- ✅ ISO 27001 certification
- ✅ Clean SOC 2 Type II

2: Perverse Incentives

What CISOs are rewarded for:

- ✅ FedRAMP authorization
- ✅ ISO 27001 certification
- ✅ Clean SOC 2 Type II

What they are *not* rewarded for:

- ❓ Breaches that never happened
- ❓ Tail-risk mitigation with no visible output
- ❓ Reducing the attack surface of legacy systems

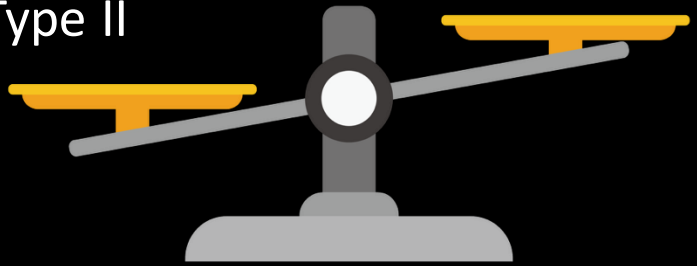

3: The Compliance Industrial Complex

Who benefits from complex compliance regimes?

- **Auditors** (Big 4 + specialty firms)
- **Consultants** (implementation, remediation)
- **Tool vendors** (*Governance, Risk & Compliance* platforms, evidence collection)
- **Certification bodies**

FedRAMP alone: 18–24 months,
$2–5M initial, $500K–$1M/year to maintain.

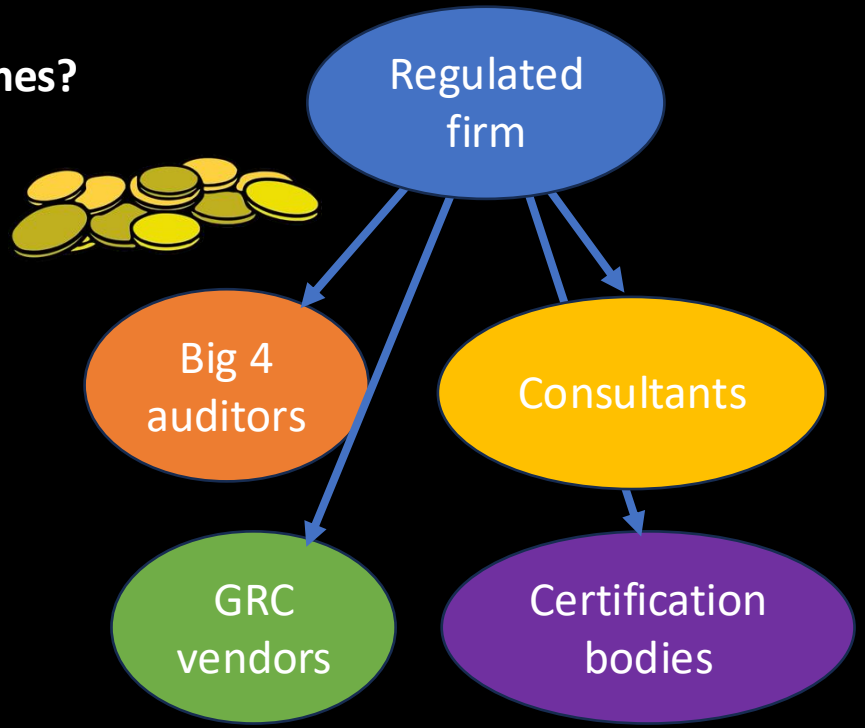

4: Regulatory Accumulation

New frameworks pile on top of old ones. Nothing gets removed.

Large orgs often maintain:

- SOC 2 Type II
- ISO 27001 / 27701
- PCI DSS
- HIPAA
- FedRAMP
- GDPR, …

Each requires separate evidence, audits, attestations.

Kelly's Proposed Alternative: A Hybrid Model

1. **Outcomes-focused** frameworks: what must be achieved, not precisely how
2. **Judgment and character** assessed alongside control verification (aviation analogy)
3. **Spring cleaning**: remove obsolete requirements
4. **Trust in expertise** within accountability

Less checklist. More professional judgment.

But Is Compliance *All* Theater?

Kelly's critique is sharp. But there's a strong counterargument:

- Post-Wyndham FTC accountability has real teeth
- Some regulation has improved measured outcomes
- Industry-informed regulation has outperformed pure checklists

Let's steelman the pro-regulation case.

Pro-Regulation Evidence: FTC Teeth

Equifax's reported breach cost:
~$1.4B total (net of insurance recoveries).

At least **half of that was the regulatory settlement**:
$700M to FTC + states + CFPB. Enabled by Wyndham, GLBA.

Without regulation, that $700M doesn't exist.

Pro-Regulation Evidence: GDPR *Did* Improve Security

1,348 UK businesses surveyed:

- Measurable improvements in:
  - **Governance**: board-level cyber responsibility up
  - **Risk management**: risk registers more widespread
  - **Data security practices**: encryption, access control
- Weaker improvements in: procurement, supply chain

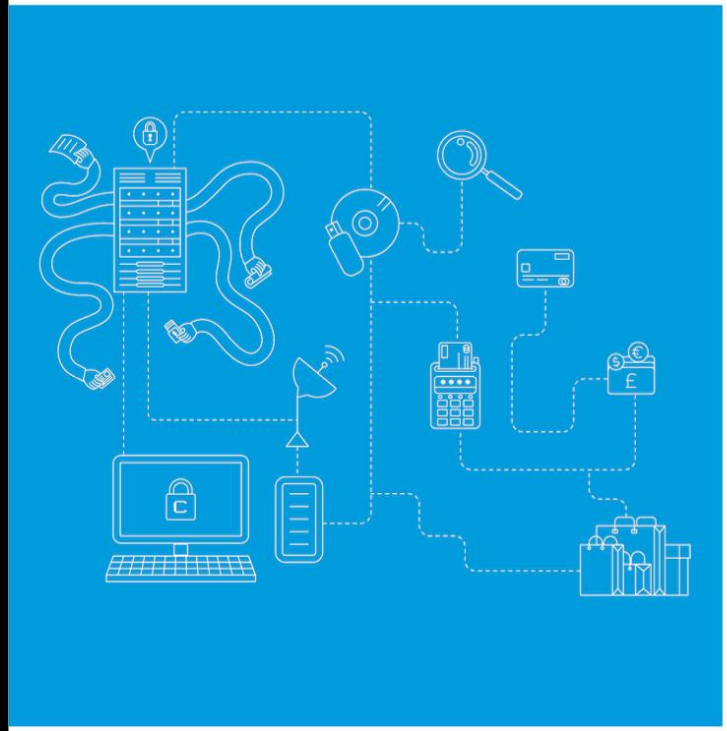

IMPACT OF THE GDPR ON CYBER SECURITY OUTCOMES

Final Report

August 2020

THE POWER OF BEING UNDERSTOOD
AUDIT | TAX | CONSULTING

RSM

Pro-Regulation Evidence: GLBA Rewritten

We saw the revised GLBA Safeguards Rule earlier.

- Replaced “reasonable” with **specific technical mandates**
- Direct response to Equifax-class failures
- Regulators can and do **learn** from big incidents

This is exactly what Kelly asks for — except it was more rules, not fewer.

FEDERAL TRADE COMMISSION

16 CFR Part 314

RIN 3084–AB35

Standards for Safeguarding Customer Information

AGENCY: Federal Trade Commission.

ACTION: Final rule.

SUMMARY: The Federal Trade Commission (“FTC” or “Commission”) is issuing a final rule (“Final Rule”) to amend the Standards for Safeguarding Customer Information (“Safeguards Rule” or “Rule”). The Final Rule contains five main modifications to the existing Rule. First, it adds provisions designed to provide covered financial institutions with more guidance on how to develop and implement specific aspects of an overall information security program, such as access controls, authentication, and encryption. Second, it adds provisions designed to improve the accountability of financial institutions' information security programs, such as by requiring periodic reports to boards of directors or governing bodies. Third, it exempts financial institutions that collect less customer information from certain requirements. Fourth, it expands the definition of “financial institution” to include entities engaged in activities the Federal Reserve Board determines to be incidental to financial activities. This change adds “finders”—companies that bring together buyers and sellers of a product or service—within the scope of the Rule. Finally, the Final Rule defines several terms and provides related examples in the Rule itself rather than incorporates them from the Privacy of Consumer Financial Information Rule (“Privacy Rule”).

DATES:

Effective date: This rule is effective January 10, 2022.

Applicability date: The provisions set forth in § 314.5 are applicable beginning December 9, 2022.

FOR FURTHER INFORMATION CONTACT: David Lincicum (202–326–2773), Katherine McCarron (202–326–2333), or Robin Wetherill (202–326–2220), Division of Privacy and Identity Protection, Bureau of Consumer Protection, Federal Trade Commission, 600 Pennsylvania Avenue NW, Washington, DC 20580.

SUPPLEMENTARY INFORMATION:

I. Background

Congress enacted the Gramm Leach Bliley Act (“GLB” or “GLBA”) in 1999.[1] The GLBA provides a framework for regulating the privacy and data security practices of a broad range of financial institutions. Among other things, the GLBA requires financial institutions to provide customers with information about the institutions' privacy practices and about their opt-out rights, and to implement security safeguards for customer information.

Subtitle A of Title V of the GLBA required the Commission and other Federal agencies to establish standards for financial institutions relating to administrative, technical, and physical safeguards for certain information.[2] Pursuant to the Act's directive, the Commission promulgated the Safeguards Rule (16 CFR part 314) in 2002. The Safeguards Rule became effective on May 23, 2003.

The current Safeguards Rule requires a financial institution to develop, implement, and maintain a comprehensive information security program that consists of the administrative, technical, and physical safeguards the financial institution uses to access, collect, distribute, process, protect, store, use, transmit, dispose of, or otherwise handle customer information.[3] The information security program must be written in one or more readily accessible parts.[4] The safeguards set forth in the program must be appropriate to the size and complexity of the financial institution, the nature and scope of its activities, and the sensitivity of any customer information at issue.[5] The safeguards must also be reasonably designed to ensure the security and confidentiality of customer information, protect against any anticipated threats or hazards to the security or integrity of the information, and protect against unauthorized access to or use of such information that could result in substantial harm or inconvenience to any customer.[6]

In order to develop, implement, and maintain its information security program, a financial institution must identify reasonably foreseeable internal and external risks to the security, confidentiality, and integrity of customer information that could result in the unauthorized disclosure, misuse, alteration, destruction, or other compromise of such information.[7] The financial institution must then design and implement safeguards to control the risks identified through the risk assessment, and must regularly test or otherwise monitor the effectiveness of the safeguards' key controls, systems, and procedures.[8] The Rule also requires the financial institution to evaluate and adjust its information security program in light of the results of this testing and monitoring, any material changes in its operations or business arrangements, or any other circumstances it knows or has reason to know may have a material impact on its information security program.[9] The financial institution must also designate an employee or employees to coordinate the information security program.[10]

Finally, the current Safeguards Rule requires financial institutions to take reasonable steps to select and retain service providers capable of maintaining appropriate safeguards for customer information and require those service providers by contract to implement and maintain such safeguards.[11]

II. Regulatory Review of the Safeguards Rule

On September 7, 2016, the Commission solicited comments on the Safeguards Rule as part of its periodic review of its rules and guides.[12] The Commission sought comment on a number of general issues, including the economic impact and benefits of the Rule; possible conflicts between the Rule and state, local, or other Federal laws or regulations; and the effect on the Rule of any technological, economic, or other industry changes. The Commission received 28 comments from individuals and entities representing a wide range of viewpoints.[13] Most commenters agreed there is a continuing need for the Rule and it benefits consumers and competition.[14]

On April 4, 2019, the Commission issued a notice of proposed rulemaking (NPRM) setting forth proposed amendments to the Safeguards Rule (the “Proposed Rule”).[15] In response, the Commission received 49 comments from various interested parties

[1] Pubic Law 106–102, 113 Stat. 1338 (1999).

[2] *See* 15 U.S.C. 6801(b), 15 U.S.C. 6805(b)(2).

[3] 16 CFR 314.2(c).

[4] 16 CFR 314.3(a).

[5] 16 CFR 314.3(a), (b).

[6] 16 CFR 314.3(a), (b).

[7] 16 CFR 314.4(b).

[8] 16 CFR 314.4(c).

[9] 16 CFR 314.4(e).

[10] 16 CFR 314.4(a).

[11] 16 CFR 314.4(d).

[12] Safeguards Rule, Request for Comment, 81 FR 61632 (Sept. 7, 2016).

[13] The 28 public comments received prior to March 15, 2019, are posted at: *https://www.ftc.gov/policy/public-comments/initiative-674.*

[14] *See, e.g.,* Mortgage Bankers Association (comment 39, NPRM); National Automobile Dealers Association (Comment 40, NPRM); Data & Marketing Association (comment 38, NPRM); Electronic Transactions Association (comment 24, NPRM); State Privacy & Security Coalition (comment 26, NPRM).

[15] FTC Notice of Proposed Rulemaking, 84 FR 13158 (April 4, 2019).

Pro-Regulation Evidence: Enlightened Capture

Thaw, "Enlightened Regulatory Capture" (*Wash. L. Rev.*, 2014):

- Compares healthcare (HIPAA) to traditionally-regulated industries
- Industry-**informed**, adaptive regulation produced practices:
  - **~4x more effective** than prescriptive checklists
- Key ingredient: **collaborative**, not just rule-imposing

Industry involvement isn't always capture-in-the-bad-sense.

PUBLIC DISCOURSE, EXPERT KNOWLEDGE, AND THE PRESS

Joseph Blocher*

Abstract: This Essay identifies and elaborates two complications raised by Robert Post's *Democracy, Expertise, and Academic Freedom*, and in doing so attempts to show how Post's theory can account for constitutional protection of the press. The first complication is a potential circularity arising from the relationships between the concepts of democratic legitimation, public discourse, and protected social practices. Democratic legitimation predicates First Amendment coverage on participation in public discourse, whose boundaries are defined as those social practices necessary for the formation of public opinion. But close examination of the relationships between these three concepts raises the question of whether public discourse and social practices can do any analytic work independent of the value of democratic legitimation, or instead are simply labels for speech that furthers it. Consideration of the press helps to illuminate the problem and a potential solution.

The second complication is the interface between expert knowledge and public discourse. Post's theory of democratic competence convincingly explains how such knowledge is created and circulated outside of public discourse. But in order to inform self-governance, expert knowledge must ultimately be disseminated *into* public discourse. The theory does not yet account for how this happens, nor how such expert knowledge can serve an informative function, given that public discourse transmutes claims of expert knowledge into statements of opinion. Again, the press serves as an illustrative and important example.

INTRODUCTION

Robert Post's *Democracy, Expertise, and Academic Freedom*[1] explains our constitutional commitment to free speech in light of two central and sometimes conflicting principles: democratic legitimation and democratic competence. In doing so, the book employs concepts that Post has carefully crafted over the past few decades, including the constitutional concept of public discourse,[2] the lexical priority of

---

* Assistant Professor, Duke Law School. Many thanks to Stuart Benjamin, Michael Gerhardt, Marin K. Levy, Robert Post, and Neil Siegel for valuable suggestions.

1. ROBERT C. POST, DEMOCRACY, EXPERTISE, AND ACADEMIC FREEDOM: A FIRST AMENDMENT JURISPRUDENCE FOR THE MODERN STATE (2012).

2. *See generally* Robert C. Post, *The Constitutional Concept of Public Discourse: Outrageous Opinion, Democratic Deliberation, and* Hustler Magazine v. Falwell, 103 HARV. L. REV. 601 (1990) [hereinafter Post, *Public Discourse*].

The Measurement Problem, Resurfaced

As you know: we can't reliably measure security ***outcomes***.

So regulation mostly measures ***inputs***:

Measurable	Not measurable
Did you patch?	Are you actually hard to breach?
Do you have a policy?	Do your controls work?
Did you pass the audit?	Is the residual risk acceptable?

The regulatory paradox: the most important goals are the least measurable. Challenges Thaw's "let experts do their thing" idea.

Part 3: The US Regulatory Toolkit

The US Regulatory Landscape

What existed in 2017:

- **GLBA Safeguards Rule** — principles-based ("reasonable" security)
- **FTC Act Section 5** — unfair/deceptive practices
- **State breach notification laws** (47 states, varying)
- **SEC securities fraud rules**

What later emerges:

- **CIRCIA (2022)** — no federal incident reporting
- **SEC cyber disclosure rule** (2023)
- **CSRB** (2022) — no post-incident learning
- **Comprehensive federal breach notification law**
- **Federal privacy law**; GDPR not yet in force

Three Categories of Approach, Post-Equifax

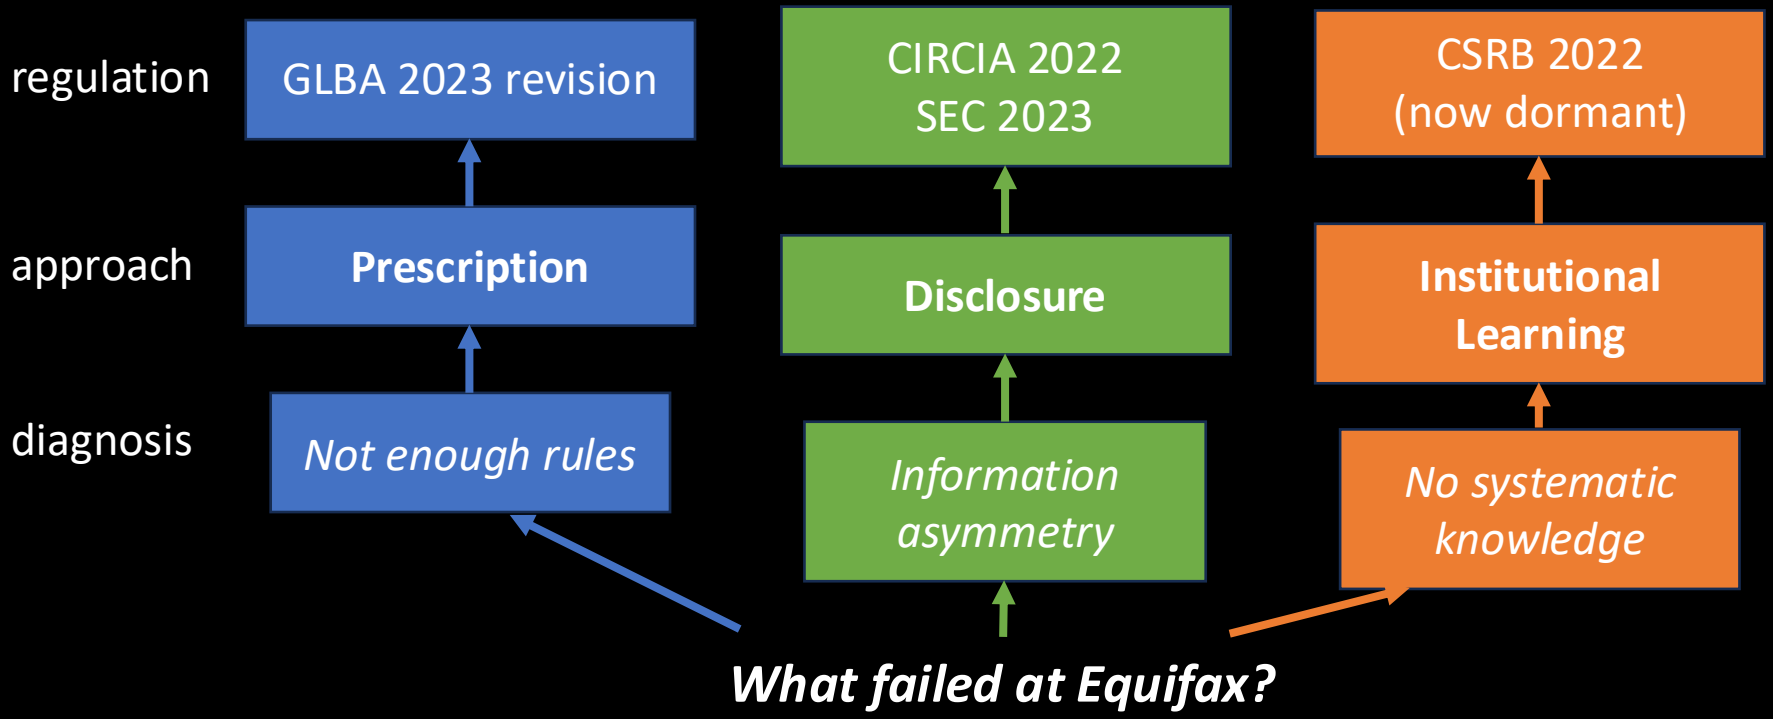

Category 1: Prescription

Specify the controls. Audit them. Enforce them.

Canonical example: **GLBA Safeguards revision (2023)** — MFA, encryption, Qualified Individual, pen testing *(covered in Part 1)*.

Two cases reveal how far prescription can actually reach in court:

- **LabMD (2018)** — vagueness of “reasonable security”
- **D-Link (2017)** — unfairness needs concrete harm

LabMD v. FTC (11th Cir., 2018)

When: complaint 2013 → FTC reversal 2016 → 11th Cir. 2018

At stake: small cancer lab; LimeWire installed on work computer exposed 9,300 patients' SSNs/insurance data via P2P

Outcome: FTC's cease-and-desist order **vacated for vagueness**

Impact: Section 5 *authority* preserved (Wyndham stands) — but "implement reasonable security" is too vague to enforce as an injunction

FTC v. D-Link (N.D. Cal., 2017)

When: complaint Jan 2017 → motion-to-dismiss ruling Sep 2017 → settlement 2019

At stake: IoT vendor marketing “EASY to secure” with hard-coded credentials, command-injection bugs, exposed signing key — **no known breach**

Outcome: unfairness claim **dismissed** (no concrete harm); deception claims survive; 20-year consent order, **no fine**

Impact: FTC must show actual or likely injury — hard to do pre-breach

Category 2: Disclosure

Force information out. Let markets, investors, and regulators price it.

Two main examples:

- **SEC Cyber Rules (2023)** — investor-facing
- **CIRCIA (2022)** — government-facing

One major case has tested (and narrowed) the limits: **SEC v. SolarWinds / Brown**.

SEC Cyber Rules (2023)

- Public companies must disclose **material** incidents within **4 business days** on Form 8-K
- Board oversight + risk-management processes in annual filings
- Conceptual shift: cyber risk as investor-relevant financial risk

Would have closed the 40-day Equifax gap

SEC v. SolarWinds / Brown — The Disclosure Limit

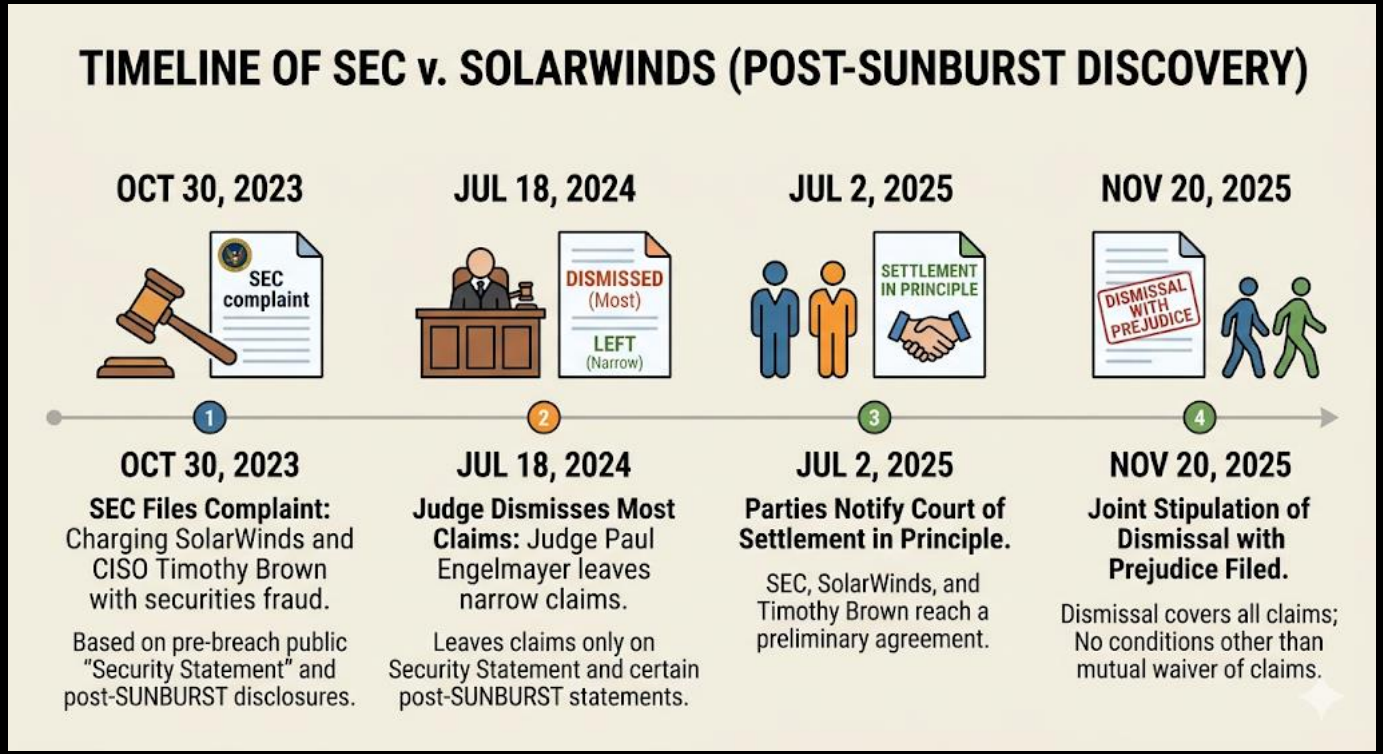

SEC v. SolarWinds / Brown — The Disclosure Limit

At stake: SEC sued SolarWinds *and* its CISO personally for securities fraud based on **pre-breach** security representations

Outcome: joint stipulation of dismissal, **with prejudice**, no settlement terms

Impact: disclosure enforcement works *prospectively* (incident reporting). Retroactive fraud claims about security statements are hard to sustain (but SEC will nevertheless consider them). First cyber-fraud case against a sitting CISO is a **defeat for the SEC**.

CIRCIA (2022)

Cyber Incident Reporting for Critical Infrastructure Act — signed March 2022.

First broad federal **incident reporting** mandate for critical infrastructure

- **72 hours** for incidents, **24 hours** for ransomware payments
- ~300,000 covered entities across 16 sectors
- Reports shielded from civil discovery — **but not anonymous**
- Final rules: delayed to **May 2026**

CIRCIA: Will It Generate the Right Data?

Aviation's learning model depends on **near-misses**, not just incidents.

CIRCIA focuses on **incidents only**.

- No institution yet analyzes the reports (CSRB now dormant)
- Reports are not anonymous — chilling effect?
- Reports go to an agency with enforcement relationships

Reporting ≠ learning. Data piled up is not data understood.

Category 3: Institutional Learning

Systematically study incidents. Publish lessons. Build cross-industry knowledge.

Cyber Safety Review Board (CSRB) — constituted 2021 → **dormant** since January 2025.

The most structurally broken of the three approaches today — and arguably the most important.

Full treatment later in the lecture

Case Law: What Have We Learned?

Case	Lesson
Wyndham (2015)	FTC *can* regulate cyber under Section 5
LabMD (2018)	Remedies must be **specific**, not “reasonable”
D-Link (2017)	Unfairness requires **concrete harm**; deception is the reliable lane
SolarWinds (2025)	CISO personal fraud liability is **hard to sustain**

Case-by-case accretion is **slow**. It also has clear limits.

How Would Reforms Have Affected the Breach?

Reform	Effect on Equifax specifically
GLBA revision (2023)	**Might have prevented**: encryption, pen testing, Qualified Individual
SEC disclosure (2023)	**Compressed the 40-day gap** to ~1 week
CIRCIA (2022)	**CISA would have had a report** — situational awareness
CSRB (2021, now dormant)	**Cross-industry learning** from a published post-incident report
GDPR (2018, EU)	Out of jurisdiction; would have lifted data-governance practices

Part 4: The EU Approach — Horizontal Regulation

The Cyber Resilience Act (Regulation 2024/2847)

First EU law imposing mandatory cybersecurity on **all products with digital elements**

- Applies regardless of manufacturer origin (market-access based)
- In force **December 2024**; reporting: **September 2026**; full compliance: **December 2027**

CRA Key Requirements

- **Secure by design** across entire product lifecycle
- Vulnerability handling and patching
- **24-hour** vulnerability reporting to national CSIRT + ENISA
- 72-hour update; 14-day final report
- **CE marking**; third-party assessment for high-risk products
- Fines up to **€15M or 2.5%** of worldwide turnover

US vs. EU: Different Philosophies

	US	EU
Approach	Sectoral + common-law	Horizontal + prescriptive
Primary lever	FTC Section 5 enforcement	Comprehensive statutes
Scope	Financial, health, etc.	Economy-wide
Flexibility	Case-by-case	Codified rules
Pace	Slow via litigation	Slow via legislation

The EU treats cybersecurity as product safety.

GDPR as an Inspirational Test Case

General Data Protection Regulation (GDPR, May 2018):

Primarily *privacy,* but with security implications.

Old enough that we have evidence.

- Has it improved security outcomes?
- Has it caused unintended harms?
- What does the track record tell us about the CRA?

GDPR: The Good

UK DCMS study (2020):

- Governance, risk management, data security practices improved
- Board-level cyber responsibility up
- Encryption and access control more widespread

Real evidence that ambitious horizontal regulation *can* shift firm behavior.

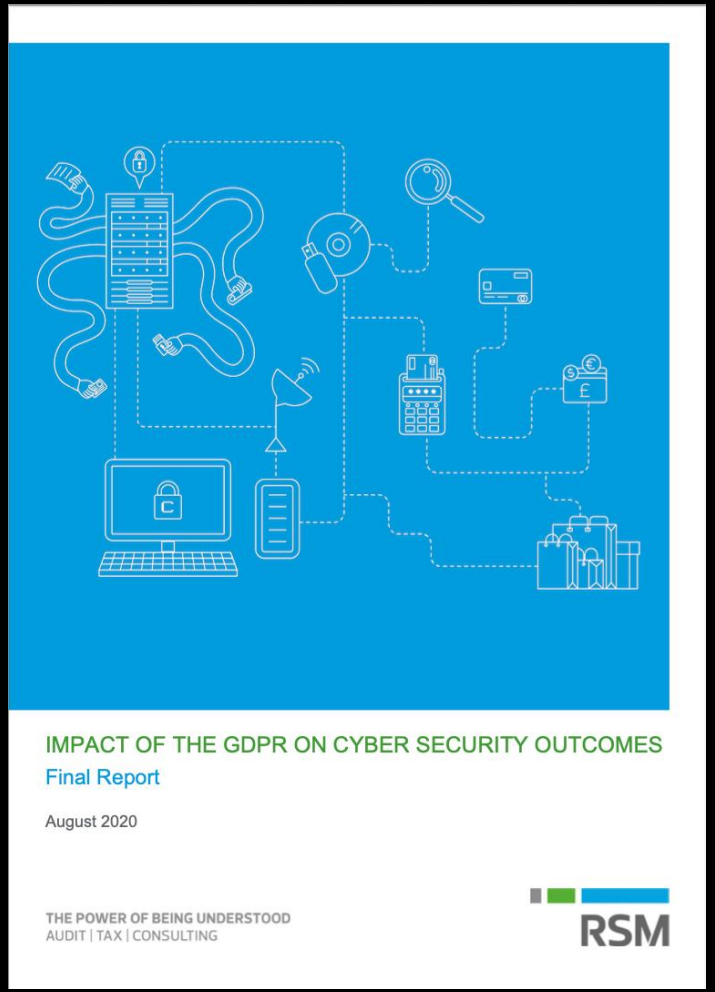

GDPR: Unintended Consequences for Security

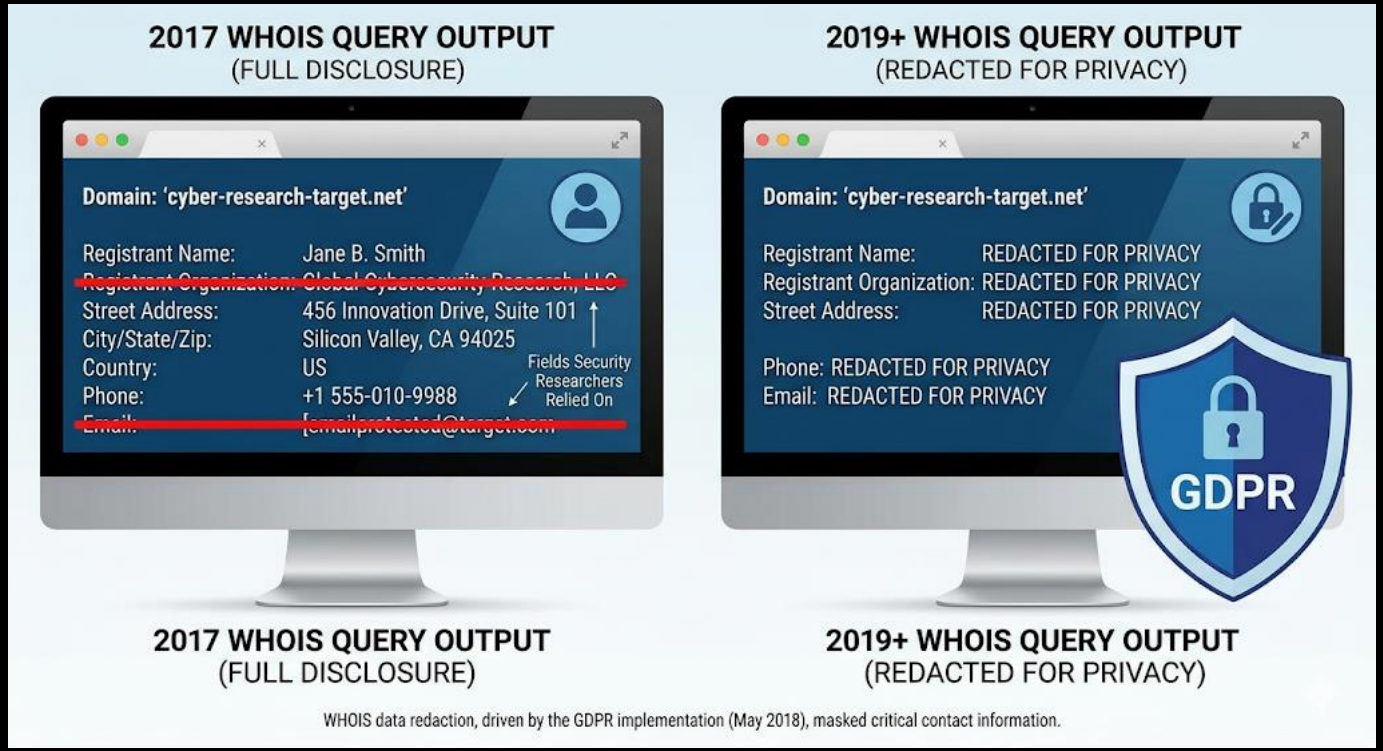

GDPR: Unintended Consequences for Security

WHOIS restrictions

- ICANN required to redact registrant data (GDPR = personal data)
- APWG (2019): domain-abuse investigations slowed dramatically
- Registrar response times: minutes → weeks
- Privacy protected — *at the expense of cybercrime investigation*

GDPR: Threat Intelligence Chilled

ENISA (2019):

- Legal uncertainty around sharing IoCs: IP addresses, email addresses = personal data?
- ISACs reported **reduced sharing volumes** initially
- Forced expensive legal review before routine intel operations

Privacy rules designed for consumer protection can block cross-border threat intel.

GDPR: Cookie Consent Theater

- Utz et al. (CCS 2020): **>90%** of users just click “accept all”
- Nouwens et al. (CHI 2020): only **11.8%** of UK cookie consent implementations met GDPR minimum requirements
- Ritual compliance, not informed choice

A literal case of compliance theater — by design.

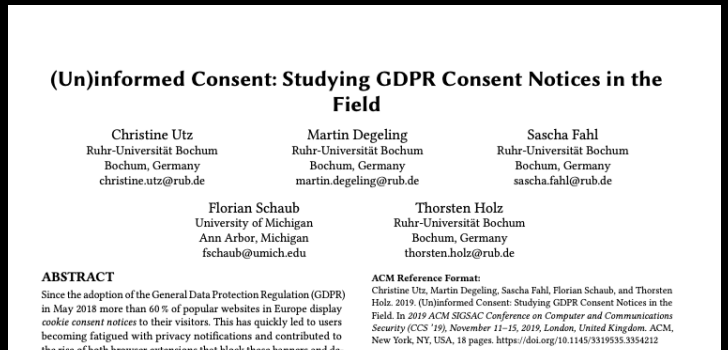

(Un)informed Consent: Studying GDPR Consent Notices in the Field

Christine Utz
Ruhr-Universität Bochum
Bochum, Germany
christine.utz@rub.de

Martin Degeling
Ruhr-Universität Bochum
Bochum, Germany
martin.degeling@rub.de

Sascha Fahl
Ruhr-Universität Bochum
Bochum, Germany
sascha.fahl@rub.de

Florian Schaub
University of Michigan
Ann Arbor, Michigan
fschaub@umich.edu

Thorsten Holz
Ruhr-Universität Bochum
Bochum, Germany
thorsten.holz@rub.de

ABSTRACT

Since the adoption of the General Data Protection Regulation (GDPR) in May 2018 more than 60 % of popular websites in Europe display *cookie consent notices* to their visitors. This has quickly led to users becoming fatigued with privacy notifications and contributed to

ACM Reference Format:
Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. 2019. (Un)informed Consent: Studying GDPR Consent Notices in the Field. In *2019 ACM SIGSAC Conference on Computer and Communications Security (CCS '19), November 11–15, 2019, London, United Kingdom.* ACM, New York, NY, USA, 18 pages. https://doi.org/10.1145/3319535.3354212

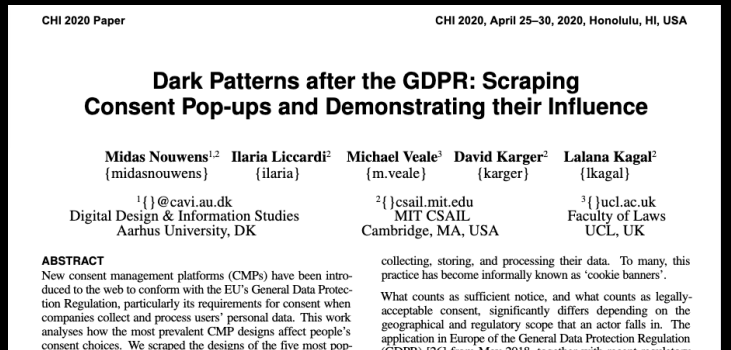

Dark Patterns after the GDPR: Scraping Consent Pop-ups and Demonstrating their Influence

Midas Nouwens[1,2] {midasnouwens}
Ilaria Liccardi[2] {ilaria}
Michael Veale[3] {m.veale}
David Karger[2] {karger}
Lalana Kagal[2] {lkagal}

[1]{}@cavi.au.dk
Digital Design & Information Studies
Aarhus University, DK

[2]{}csail.mit.edu
MIT CSAIL
Cambridge, MA, USA

[3]{}ucl.ac.uk
Faculty of Laws
UCL, UK

ABSTRACT

New consent management platforms (CMPs) have been introduced to the web to conform with the EU's General Data Protection Regulation, particularly its requirements for consent when companies collect and process users' personal data. This work analyses how the most prevalent CMP designs affect people's consent choices. We scraped the designs of the five most pop-

collecting, storing, and processing their data. To many, this practice has become informally known as 'cookie banners'.

What counts as sufficient notice, and what counts as legally-acceptable consent, significantly differs depending on the geographical and regulatory scope that an actor falls in. The application in Europe of the General Data Protection Regulation

Post-GDPR: Compliance Costs and Competition

- Goldberg, Johnson & Shriver (2022)
- Johnson, Shriver & Goldberg (2022)
- Jia, Jin & Wagman (2020)

Enforcement concentrated in France (CNIL), Ireland (DPC); Ireland criticized for slow Big Tech action

Regulation has real costs --- and can entrench incumbents.

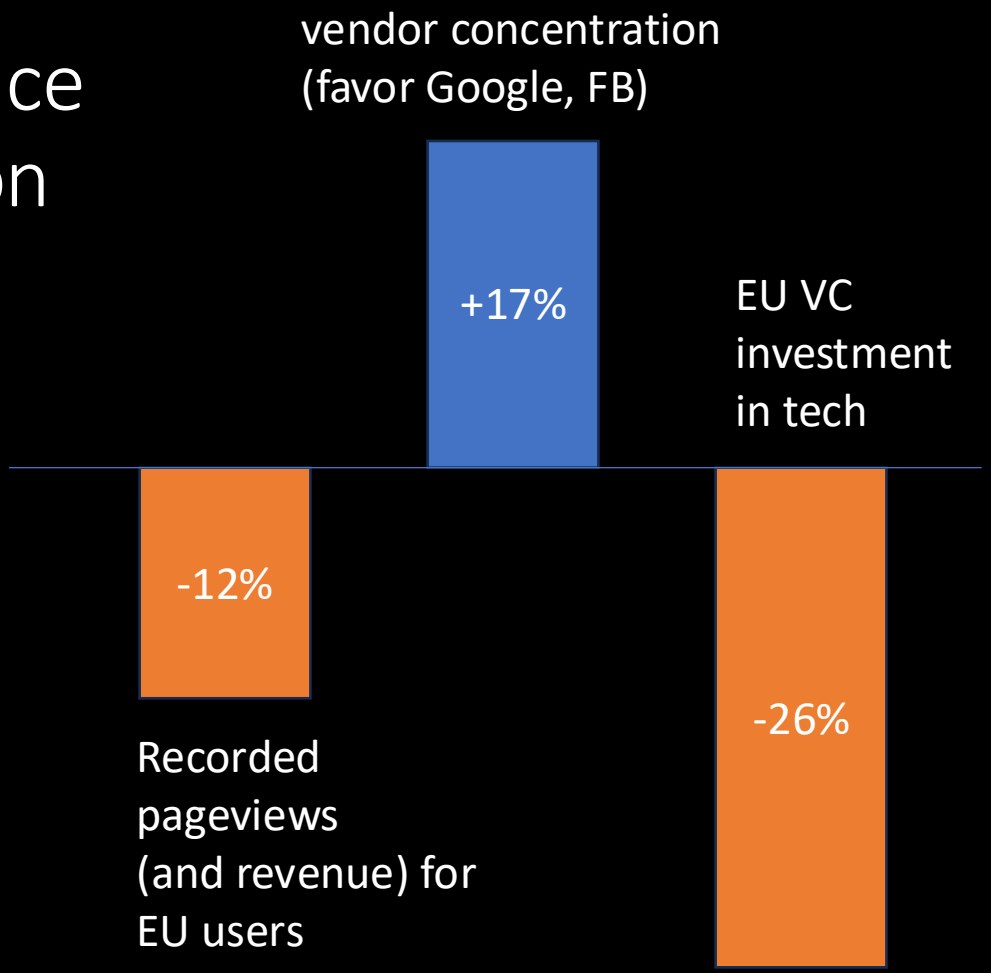

Lesson for CRA

GDPR shows horizontal regulation:

- ✅ Can shift firm behavior (governance, practices)
- ⚠️ Creates unintended security harms (WHOIS, threat intel)
- ⚠️ Produces compliance theater at scale (consent banners)
- ⚠️ Imposes costs unevenly (SMEs, incumbents)
- ⚠️ Enforcement concentrates (and fails unevenly)

CRA's scope is larger. Each of these risks is larger too.

Back to CRA: What It Doesn’t Do Well

- “Secure by design” is **vague** until harmonized standards exist
- Technical standards (CEN, CENELEC, ETSI) **still being drafted** (slated for Oct 2026)
- Monitoring every connected product is **impossible**
- Expected pattern: **enforcement-by-exception** after big incidents — like GDPR

A law without its standards is a promise, not a requirement.

CRA: The Open Source Controversy

- ~76% of most software products is **open source dependencies**
- First draft: OSS maintainers liable for commercial use of their code
- Community pushback → revised:
  - Non-profit-motive OSS exempted
  - New “open source steward” category for foundations

Still unresolved: what happens when a critical library’s maintainer can’t comply?

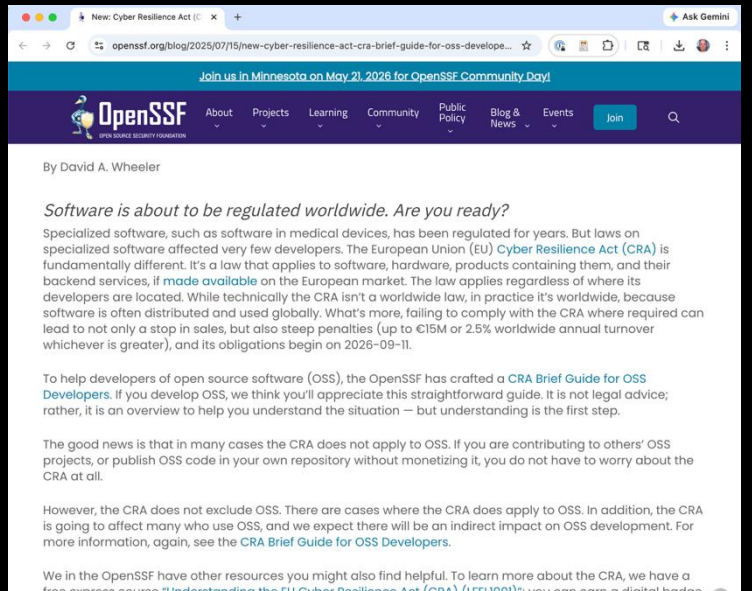

The Cybersecurity Act 2 (CSA2)

New element added 2026: **high-risk supplier** designation for countries posing **non-technical risks**.

- Laws requiring vulnerability disclosure to that country's authorities
- Absence of judicial or democratic oversight
- Credible indications of state-sponsored cyber activity

NIS2 entities prohibited from using components from designated suppliers. Phase-out required.

Fines up to **7%** of worldwide turnover.

CSA2: Cyber Regulation or Industrial Policy?

Huawei's objection:

> Restricting suppliers by country of origin rather than technical evidence violates EU principles of fairness, non-discrimination, and proportionality.

A real tension:

- Could purely technical assessment reach the same conclusions? Maybe.
- CSA2 *explicitly* chose a non-technical basis. That's notable.

Will Industry Pushback Weaken CRA/CSA2?

US precedent: **FCC Anti-Bot Code of Conduct (2013)**

- Attempted mandatory cybersecurity obligations on ISPs
- Industry resistance: “we can’t kick customers offline”
- Ended up **voluntary → toothless**

Similar dynamics now:

- Telecom lobby: CSA2 proposals = “billions of euros” in costs
- OSS community already forced CRA revisions

Regulation: US vs. EU Divergence

US lighter-touch:

- **Cyber Trust Mark** (voluntary labeling, 2024)
- **CIRCIA** (reporting for critical infra)
- No CRA equivalent

Open question: will EU horizontal regulation outperform US sectoral enforcement?

Thaw's 2014 Idea: Blended Regulation

	No concrete controls	Specifies concrete controls
No industry input	(rare in cyber)	State SBNs *(directive)*
Industry co-develops	HIPAA, GLBA (original), NIST CSF *(MBRD)*	**GLBA 2023 revision** *(blended — Thaw's ideal)*

Empirical finding: **Management-Based Regulatory Delegation (MBRD) > directive alone**,

but a ***blend*** **of the two beats either alone**.

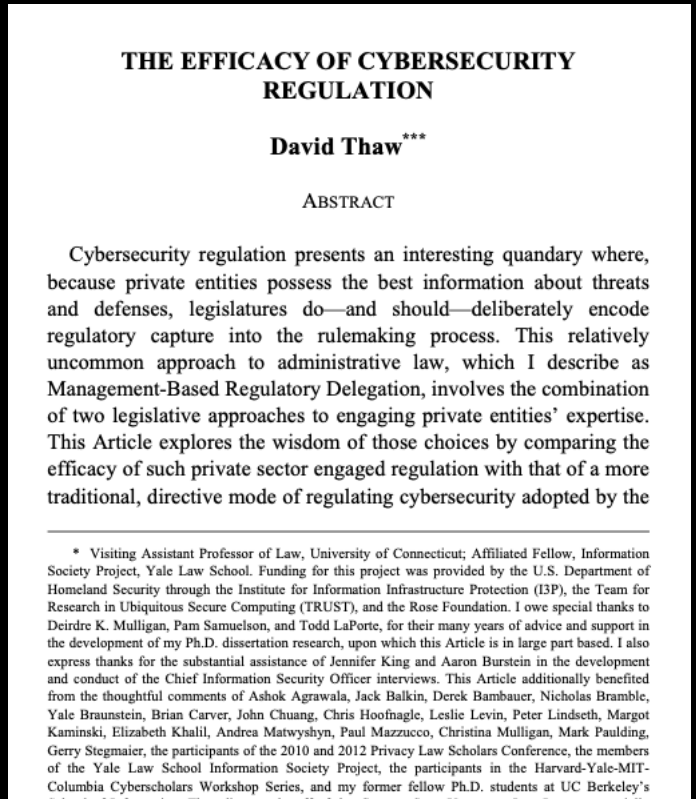

THE EFFICACY OF CYBERSECURITY REGULATION

David Thaw***

ABSTRACT

Cybersecurity regulation presents an interesting quandary where, because private entities possess the best information about threats and defenses, legislatures do—and should—deliberately encode regulatory capture into the rulemaking process. This relatively uncommon approach to administrative law, which I describe as Management-Based Regulatory Delegation, involves the combination of two legislative approaches to engaging private entities' expertise. This Article explores the wisdom of those choices by comparing the efficacy of such private sector engaged regulation with that of a more traditional, directive mode of regulating cybersecurity adopted by the

* Visiting Assistant Professor of Law, University of Connecticut; Affiliated Fellow, Information Society Project, Yale Law School. Funding for this project was provided by the U.S. Department of Homeland Security through the Institute for Information Infrastructure Protection (I3P), the Team for Research in Ubiquitous Secure Computing (TRUST), and the Rose Foundation. I owe special thanks to Deirdre K. Mulligan, Pam Samuelson, and Todd LaPorte, for their many years of advice and support in the development of my Ph.D. dissertation research, upon which this Article is in large part based. I also express thanks for the substantial assistance of Jennifer King and Aaron Burstein in the development and conduct of the Chief Information Security Officer interviews. This Article additionally benefited from the thoughtful comments of Ashok Agrawala, Jack Balkin, Derek Bambauer, Nicholas Bramble, Yale Braunstein, Brian Carver, John Chuang, Chris Hoofnagle, Leslie Levin, Peter Lindseth, Margot Kaminski, Elizabeth Khalil, Andrea Matwyshyn, Paul Mazzucco, Christina Mulligan, Mark Paulding, Gerry Stegmaier, the participants of the 2010 and 2012 Privacy Law Scholars Conference, the members of the Yale Law School Information Society Project, the participants in the Harvard-Yale-MIT-Columbia Cyberscholars Workshop Series, and my former fellow Ph.D. students at UC Berkeley's

Twelve Years On: Thaw Partly Vindicated

✅ **“Reasonableness impedes compliance”**: LabMD (2018) — 11th Cir. vacated FTC order precisely because “reasonable” is unenforceable. Exactly Thaw’s prediction.

✅ **Blended regulation works**: 2023 GLBA revision adds specific directive mandates (MFA, encryption, Qualified Individual) **on top of** the existing MBRD structure. This is Thaw’s prescription, enacted.

✅ **MBRD model expanded**: NIST CSF 2.0 (2024), voluntary and industry-co-developed, is pure MBRD in Thaw’s sense.

Twelve Years On: Partly Overrun

⚠ **Horizontal regulation**: GDPR and CRA bet against Thaw — economy-wide, prescriptive, limited industry-delegation. Evidence is mixed: DCMS shows governance uplift, but cookie theater, WHOIS harms, SME burdens support his caution.

⚠ **New disclosure mandates** don't fit cleanly: SEC 2023 and CIRCIA are information-forcing for *different audiences* (investors, government) — not the state SBN consumer-notification he studied.

⚠ **Personal liability for executives**: Drizly CEO, SolarWinds/Brown. Thaw's entity-level framework doesn't speak to this.

Part 5: Incident Reporting and the Cyber NTSB

Reporting & Investigation: Public Health Analogy

Disease reporting → epidemiology → population-level interventions.

Could cyber incident reporting do the same?

Four principles from public health

1. Expert governance
2. Reporting minimization
3. De-identification
4. Use limitations

None of the current regimes (CIRCIA, SEC, CRA, NIS2) fully implement these principles

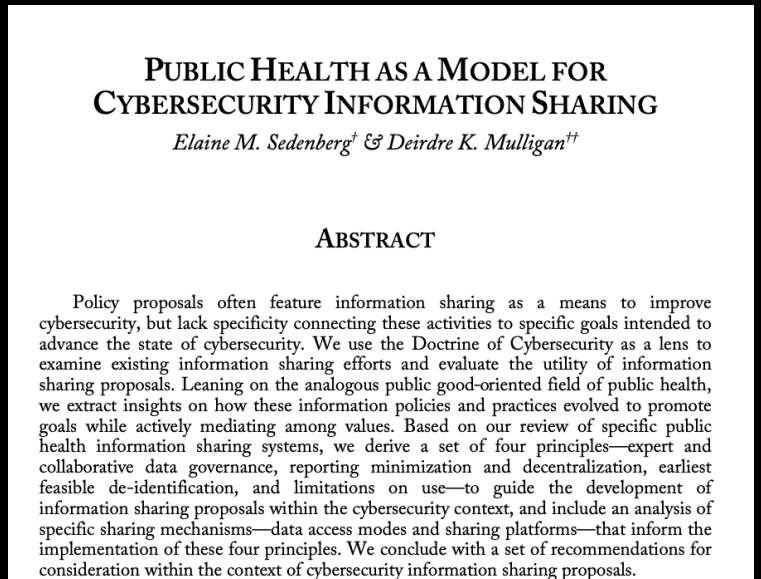

PUBLIC HEALTH AS A MODEL FOR CYBERSECURITY INFORMATION SHARING

Elaine M. Sedenberg† & Deirdre K. Mulligan††

ABSTRACT

Policy proposals often feature information sharing as a means to improve cybersecurity, but lack specificity connecting these activities to specific goals intended to advance the state of cybersecurity. We use the Doctrine of Cybersecurity as a lens to examine existing information sharing efforts and evaluate the utility of information sharing proposals. Leaning on the analogous public good-oriented field of public health, we extract insights on how these information policies and practices evolved to promote goals while actively mediating among values. Based on our review of specific public health information sharing systems, we derive a set of four principles—expert and collaborative data governance, reporting minimization and decentralization, earliest feasible de-identification, and limitations on use—to guide the development of information sharing proposals within the cybersecurity context, and include an analysis of specific sharing mechanisms—data access modes and sharing platforms—that inform the implementation of these four principles. We conclude with a set of recommendations for consideration within the context of cybersecurity information sharing proposals.

The Cyber Safety Review Board (CSRB)

Constituted **2021** by executive order.
Modeled on the NTSB.

Investigated:

- **Lapsus$** group tactics
- **Microsoft Exchange / Storm-0558** intrusion

CSRB: *Dormant* Since January 2025

CSRB: *Dormant* Since January 2025

- January 2025: Trump administration dismissed **all advisory committee members** at DHS, including CSRB
- Halted ongoing investigation of **Salt Typhoon** (Chinese intrusions into US telecom)
- Executive order establishing CSRB remains in place
- DHS nominee Troy Edgar: will “reconstitute at the right time”
- As of early 2026: **no replacements**

CSRB: Structural Weaknesses

Even at its best, CSRB had none of what makes NTSB work.

Authority	NTSB	CSRB
Subpoena power	✓	X
Participant legal immunity (49 USC §1154)	✓	X
Independence from enforcement agency	✓	X
Statutory foundation	✓	X (exec. order)
Anonymous near-miss channel (ASRS)	✓	X
Reports barred from litigation use	✓	X

What Would a Real Cyber NTSB Need?

Workshop of 70+ experts:

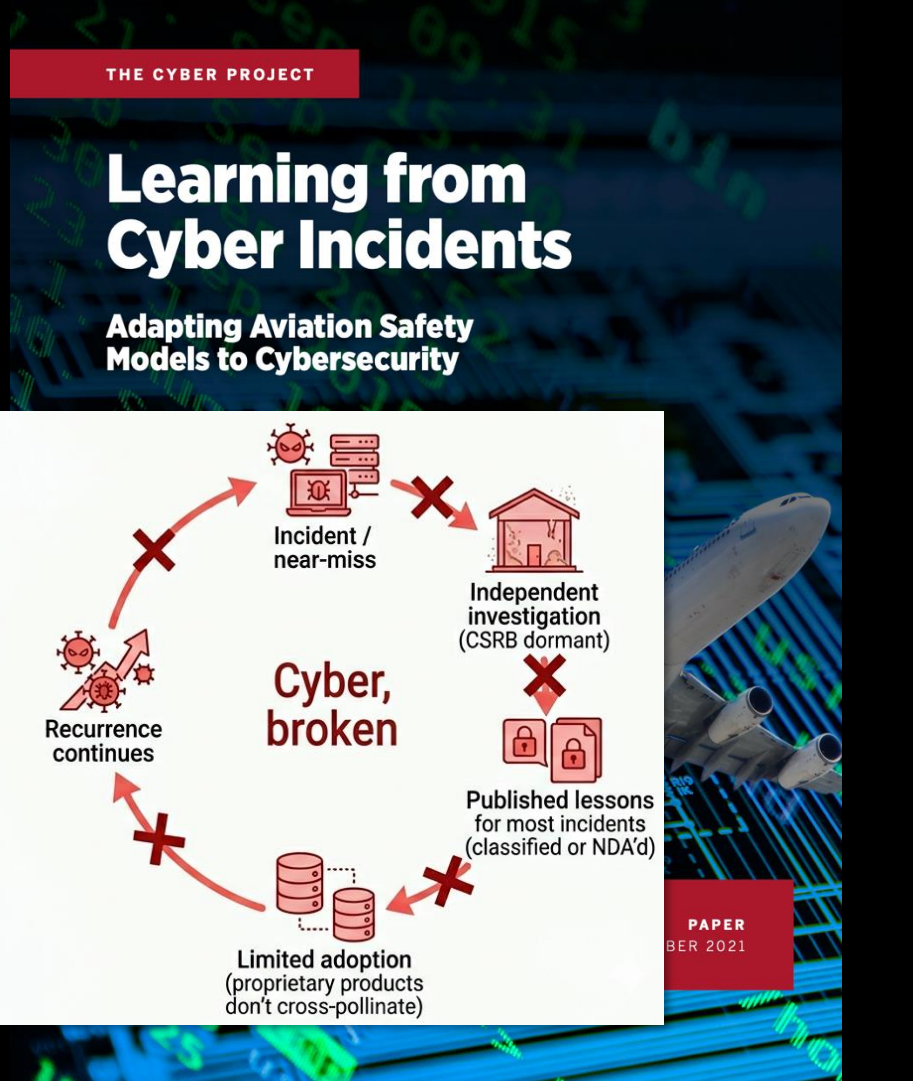

What Would a Real Cyber NTSB Need?

- **Independent** investigation body
- **Authority to compel** cooperation (subpoena)
- **Non-punitive** reporting protections (like ASRS)
- **Anonymous** near-miss channel
- **Institutional memory** and systematic analysis
- **Statutory** foundation, not executive order

Closing the Loop

Last time:

> “Cybersecurity lacks the institutional learning infrastructure that aviation has.”

This time:

> “Regulation alone won’t build it — regulation devolves into process compliance. Institutional learning requires **different** institutions.”

A Third Mechanism

We've now seen two:

- **Risk-based investment**: data problems, measurement problems
- **Regulation and compliance**: compliance ≠ security, process-heavy

Next lecture: can insurance markets price what regulators can't measure?

References

Equifax and FTC enforcement:

- FTC v. Wyndham Worldwide, 799 F.3d 236 (3d Cir. 2015)
- LabMD, Inc. v. FTC, 894 F.3d 1221 (11th Cir. 2018)
- SEC v. SolarWinds, No. 1:23-cv-09518 (S.D.N.Y.)
- FTC v. D-Link, No. 3:17-cv-00039 (N.D. Cal.)
- US-GAO, "Actions Taken by Equifax and Federal Agencies," GAO-18-559

Compliance critique / defense:

- Kelly, "The Compliance Theatre," CSAN 2025
- Marotta & Madnick, "Compliance / Cybersecurity Interplay," MIT CISL 2020
- Thaw, "Enlightened Regulatory Capture," *Wash. L. Rev.* 2014
- Thaw, "The Efficacy of Cybersecurity Regulation," *Ga. St. U. L. Rev.* 2014

EU regulation:

- Cyber Resilience Act (Regulation 2024/2847)
- EU Commission CSA2 Proposal, January 2026
- UK DCMS, "Impact of GDPR on Cyber Security Outcomes," 2020
- Nouwens et al., "Dark Patterns after the GDPR," CHI 2020
- Utz et al., "(Un)informed Consent," CCS 2020

Incident reporting / learning:

- CIRCIA (2022)
- SEC Cybersecurity Rules, 17 CFR 229.106 (2023)
- Knake, Shostack & Wheeler, "Learning from Cyber Incidents," Belfer 2021
- Sedenberg & Mulligan, "Public Health as a Model…," Berkeley Tech. L.J. 2016
- Romanosky, Telang & Acquisti, *JPAM* 2011
- Greenwood & Vaaler, *Rev. L. & Econ.* 2023
- Tannenbaum & Rudawski, "SolarWinds Dismissed," Harvard 2025
- Wolff, "Harmonizing U.S. Cybersecurity Regulations," SSRN 2024
- GAO-25-108436, "Cybersecurity Regulations: Industry Perspectives," 2025

Appendix: AI Regulation — Additional Detail

A New Frontier (Same Debates)

Earlier lecture: **GenAI compresses attacker economics.**

- CVE-Bench: LLMs solve **13%** of critical CVEs
- AIxCC finals: **86%** vulnerability ID on 54M LoC at $152/task
- Anthropic: **500+** zero-days in well-audited open source
- Hong Kong 2024 deepfake CFO fraud: **$25M**

Regulators are responding. Every debate from this lecture returns:

- Horizontal vs. sectoral — Prescription vs. disclosure vs. learning
- Measurement problem — Compliance theater

EU AI Act (Regulation 2024/1689)

Adopted June 2024; entered into force Aug 2024. Staged rollout:

- **Feb 2025**: prohibited practices (social scoring, untargeted biometrics)
- **Aug 2025**: governance + general-purpose AI (GPAI) obligations
- **Aug 2026 → Dec 2027** *(delayed via Digital Omnibus, April 2026)*: high-risk systems

Risk-tiered horizontal regulation: prohibited / high-risk / limited-risk / minimal-risk

Fines up to **€35M or 7%** worldwide turnover.

Same architecture as GDPR and CRA.

AI Act: Cyber-Relevant Provisions

GPAI with "systemic risk" (presumed at $>10^{25}$ FLOPs, Art. 51) owes under **Art. 55**:

- Adversarial testing / **red-teaming**
- **"Adequate cybersecurity"** for model + physical infrastructure
- **Serious incident reporting** to the EU AI Office
- Model evaluations, risk mitigation, documentation

Compliance path: the **GPAI Code of Practice** (July 2025, voluntary).

Notable: **Meta refused to sign**. Google, OpenAI signed with carve-outs.

US: Regulatory Whiplash

- **Oct 2023**: Biden EO 14110 — 10^{26} FLOP reporting, mandatory red-teaming, NIST AI Safety Institute
- **Jan 2025**: Trump **rescinded** EO 14110 (EO 14179 “Removing Barriers”)
- **June 2025**: AI Safety Institute → **CAISI** (Center for AI *Standards & Innovation*) — mission shifted from pre-deployment safety to “standards”
- **May 2025**: BIS **AI Diffusion Rule rescinded** (model-weight + chip export controls)
- **Dec 2025**: FY2026 NDAA — DoD-specific AI requirements; DeepSeek/PRC model ban
- Congress **dropped** proposed moratorium on state AI laws

Federal pre-deployment AI regulation is effectively paused. States and sectoral regulators fill the gap.

California SB 53: Disclosure, Not Prescription

Sept 2024: Newsom **vetoed SB 1047** — would have mandated safety plans + kill switches. Veto rationale: compute thresholds give "false sense of security."

Sept 2025: Signed **SB 53**. Most provisions effective **Jan 1, 2026**.

Disclosure-only regime for frontier labs (>10^{26} FLOPs):

- Annual public **safety framework**
- Pre-deployment **transparency reports**
- **Critical incident reporting** to CalOES: **15 days; 24h for imminent danger**
- Whistleblower protections; civil penalties up to **$1M / violation**

~5–8 labs in scope. **This is Category 2 — disclosure, not prescription.**

The Same Debates Return

Debate	Cyber regulation	AI regulation
Horizontal vs. sectoral	EU CRA vs. US FTC §5	EU AI Act vs. US state-by-state
Prescription vs. disclosure	GLBA vs. SEC 8-K	EU Art. 55 vs. CA SB 53
Measurement problem	Can't measure security	Can't measure safety; FLOPs are a proxy
Compliance theater	SOC 2 / PCI-DSS	GPAI Code self-declaration
Institutional learning	CSRB → dormant Jan 2025	**DHS AI Safety Board → dormant Jan 2025**

Same pathologies. Moving faster.

AI + Cyber: The Immediate Stakes

Regime	Reaches AI-as-cyber-weapon?
EU AI Act, Art. 55	**Yes** — systemic-risk GPAI must do adversarial testing
US federal	No, except chip export controls + DoD
CA SB 53	Indirectly, via "critical safety incident" reporting
Product liability (CrowdStrike-style)	No

Who regulates AI systems that *are themselves* offensive cyber tools?

Dual-use foundation models: no jurisdiction owns this.