Law, Human Subjects, Ethics and other limitations on research acceptability

Stefan Savage
University of California, San Diego

Questions?

What is all of this about?

- Research and how it is conducted is rarely (ever?) value neutral
- Some potential issues with research methodology or outcomes
  - Violate laws
  - Causes harm
  - Violate norms
  - Conflict with strongly held beliefs
- Impacts whether research is:
  - Pursued (or not) or funded (or not)
  - Approved internally (e.g., via IRB)
  - Accepted for publication (program committees)
  - The subject of public ire/scrutiny or lawsuits

Some historical context in the US

- Prior to 1906 no regulations on use of human subjects in research; Pure Food and Drug Act
- Nuremberg Code (1948)
  - "The voluntary consent of the human subject is absolutely essential" (no force of law)
- Thalidomide (late 1950s)
  - 1962 Kefauver Amendments to Food and Drug Act
- Tuskegee Syphilis Study (1932-1972)
  - Study stopped when became public; President Clinton apologizes in 1997
- National Research Act (1974)
  - National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research – charged with establishing basic ethical principles for human subjects research – Belmont Report (1979)
- HHS issues first version of "common rule" in 1981; adopted by most other depts that fund research in 1991 (FDA has some similar, but different rules)
  - 45 CFR Part 46
- Creates institutional obligations on all research that is Federally funded by those departments who have formally adopted the common rule (17 Depts; most but not all)

Some historical context in computer science

- Circa 2008, the CS measurement/security community not significantly engaged in these questions
  - Ethics, IRB, legal review rare
  - Routinely sniffing full content on ISP links, etc.

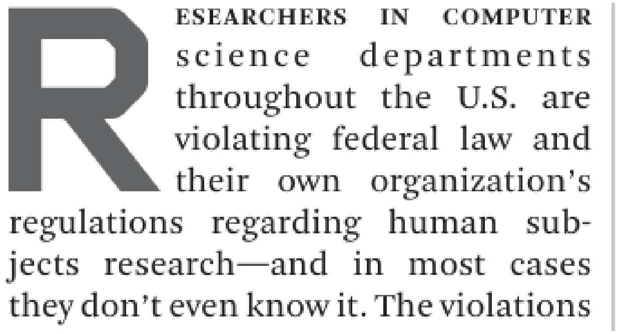

RESEARCHERS IN COMPUTER science departments throughout the U.S. are violating federal law and their own organization's regulations regarding human subjects research—and in most cases they don't even know it. The violations

Viewpoint article in ACM, June 2010

- A handful of papers (all security) caused people to notice this:
  - Designing and Conducting Phishing Experiments, 2007
  - Shining Light in Dark Places: Understanding the Tor Network, 2008
  - Spamalytics, 2008
  - Learning More about the Underground Economy: A Case-Study of Keyloggers and Dropzones, 2008
  - Your Botnet is My Botnet: Analysis of a Botnet Takeover, 2009

Speaking of Spamalytics…

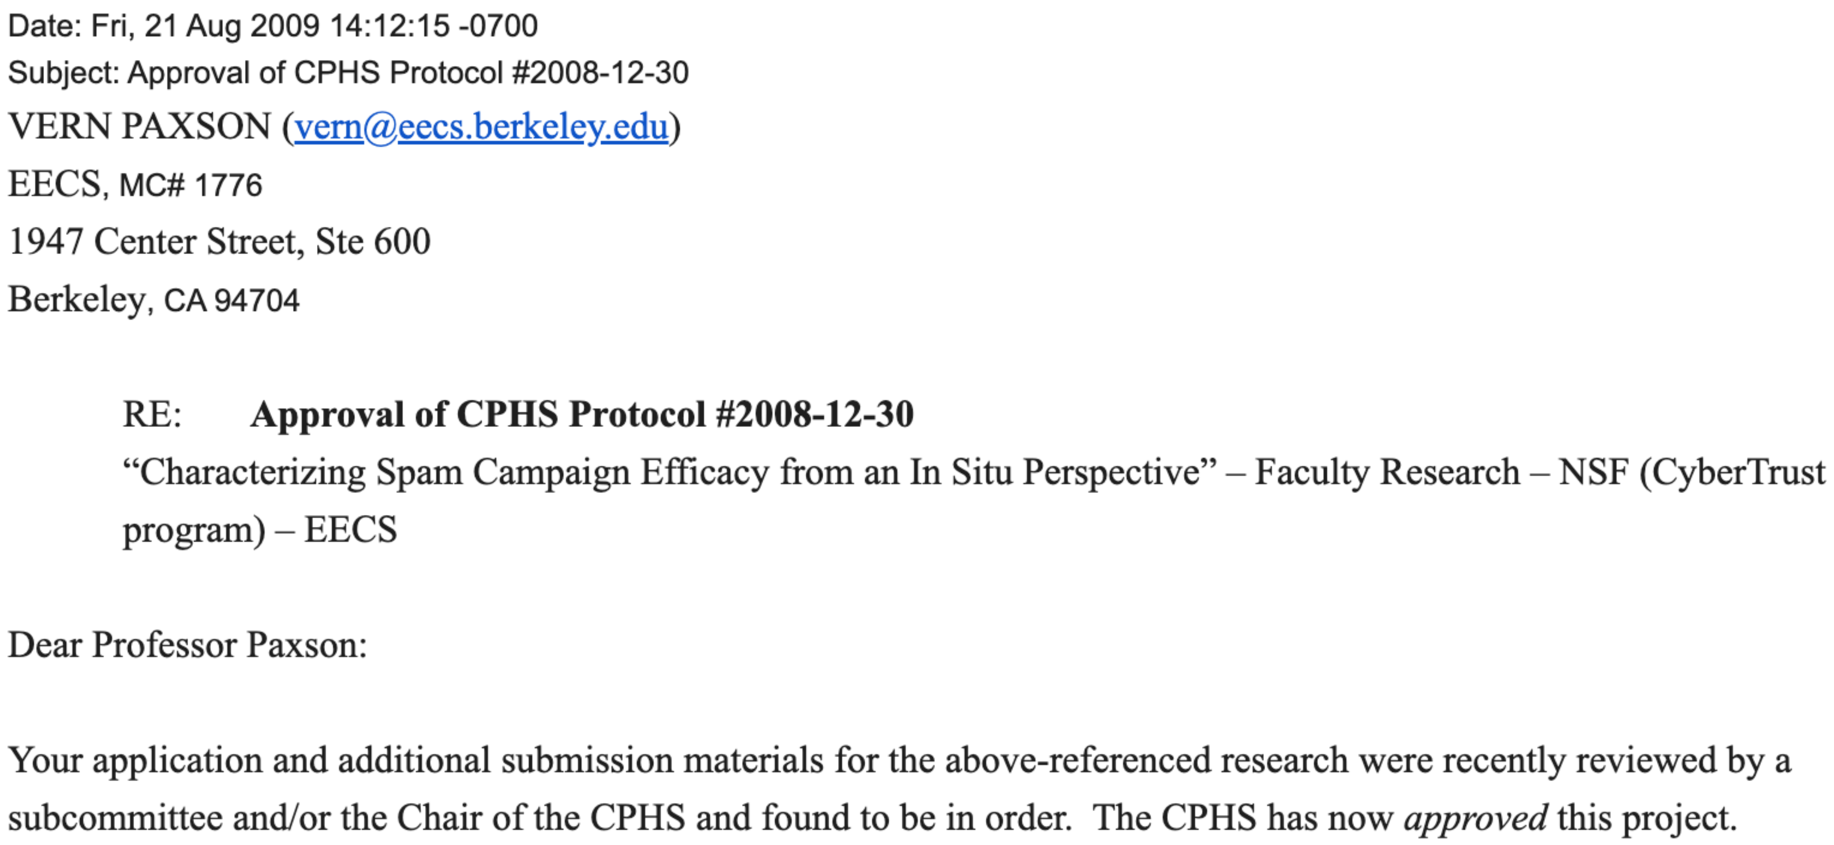

Date: Fri, 21 Aug 2009 14:12:15 -0700
Subject: Approval of CPHS Protocol #2008-12-30

VERN PAXSON (vern@eecs.berkeley.edu)
EECS, MC# 1776
1947 Center Street, Ste 600
Berkeley, CA 94704

RE: **Approval of CPHS Protocol #2008-12-30**
"Characterizing Spam Campaign Efficacy from an In Situ Perspective" – Faculty Research – NSF (CyberTrust program) – EECS

Dear Professor Paxson:

Your application and additional submission materials for the above-referenced research were recently reviewed by a subcommittee and/or the Chair of the CPHS and found to be in order. The CPHS has now *approved* this project.

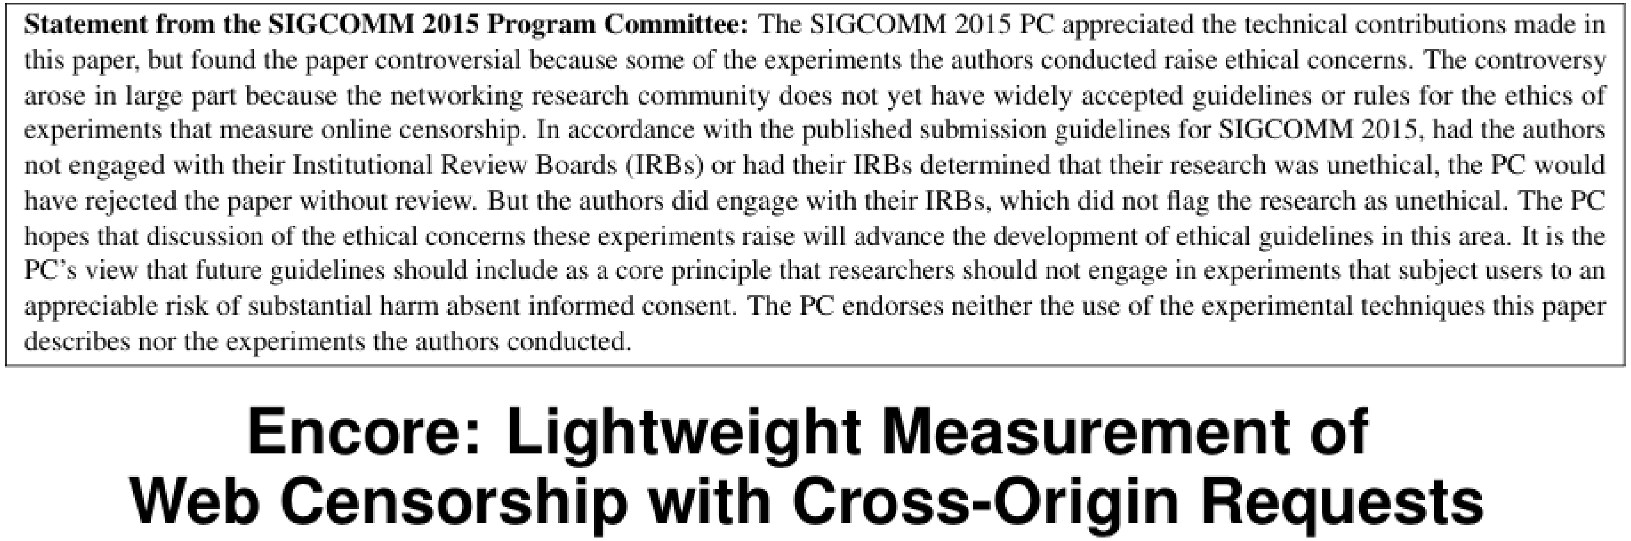

Statement from the SIGCOMM 2015 Program Committee: The SIGCOMM 2015 PC appreciated the technical contributions made in this paper, but found the paper controversial because some of the experiments the authors conducted raise ethical concerns. The controversy arose in large part because the networking research community does not yet have widely accepted guidelines or rules for the ethics of experiments that measure online censorship. In accordance with the published submission guidelines for SIGCOMM 2015, had the authors not engaged with their Institutional Review Boards (IRBs) or had their IRBs determined that their research was unethical, the PC would have rejected the paper without review. But the authors did engage with their IRBs, which did not flag the research as unethical. The PC hopes that discussion of the ethical concerns these experiments raise will advance the development of ethical guidelines in this area. It is the PC's view that future guidelines should include as a core principle that researchers should not engage in experiments that subject users to an appreciable risk of substantial harm absent informed consent. The PC endorses neither the use of the experimental techniques this paper describes nor the experiments the authors conducted.

Encore: Lightweight Measurement of Web Censorship with Cross-Origin Requests

- Eventual internalization of PC responsibility for ethics assessment beyond IRB
- Expectations of Ethics sections for papers with potentially controversial methodologies
  - This last USENIX security required an Ethics section from **all papers**
- Creation of Ethics Review Committees to specialize in such decisions

Potential issues for research acceptability

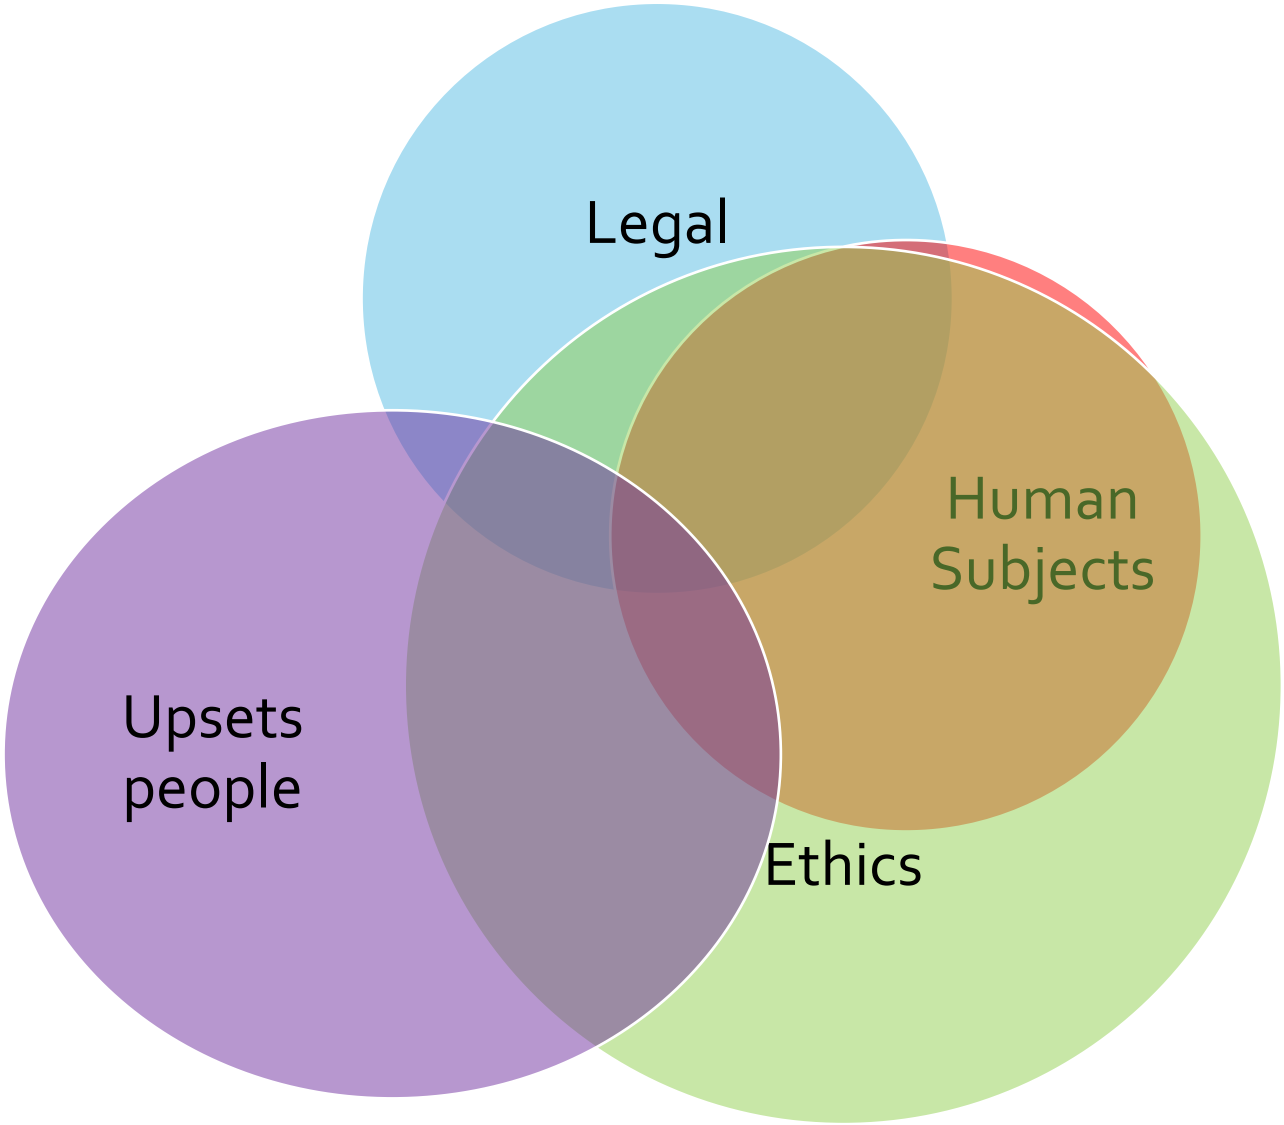

Legal

- This should be the easy one – can we agree that we shouldn’t do research that involves doing illegal things?
- Illegal how? What does it mean to be illegal?

Civil law (i.e., disputes between people and/or corps)

- Typically:
  - Contracts (violations of an agreement between consenting parties)
  - Torts (civil wrongs against people or property) – e.g., negligence, liability, interference, etc
  - Remedy: payment, return of property, stopping behavior, etc.
- But who decides if you broke the law?
  - Well, other party sues you and then you go to court… judge or jury decides
- So you don’t know ahead of time?
  - Well, many times its very clear, but sometimes its not – that’s why we need courts.

Some ”hypothetical” examples (re: civil law)

- Some companies have sued over web crawlers – is web crawling illegal?
  - Aside: not under CFAA (Van Buren decision); but other theories still exist (tortious interference, trespass to chattels, etc. but tricky… what are damages?)
  - In the US, we *generally* treat pure crawling as legal, but can get more murky if something you do might interferes with normal use of site (see “red lines” paper)
- You’re doing research into the security of voting machines and you get a “cease and desist” letter form the company. Is the research illegal then?

Some ”hypothetical” examples (re: civil law)

- You violate the “acceptable use policy” (AUP) of a Web site. Illegal?
  - (e.g., you create fake accounts on a employment site, in violation of their terms, to test if otherwise indistinguishable applications get fewer interviews if the applicant identifies themselves as being of a racial minority)
- You reverse engineer a piece of software and find a vulnerability?
  - You publish a paper about the bug without disclosing it to the vendor, they lose customers and sue for tortious interference. Illegal?
  - What if there was a “click-wrap” contract that forbade reverse engineering?
  - What if you use a pirated copy?

Quick aside: how lawyers talk about this stuff

- Good luck getting a lawyer to tell you that something ambiguous is legal or illegal. That’s not what they do.
- They will talk about the risk of certain actions and their comfort with the risk
- Ranges of things that can impact practical risk for civil litigation:
  - You’re an academic pursuing truth and science; no financial benefits?
  - Are there clear damages to some party?
  - Did you act in good faith to minimize potential risks?
  - Did you “borrow trouble” by taking nasty adversarial tone with potential adverse party?
  - Did you do due diligence with legal counsel to show that you care about such things?

Ok, civil law seems tricky… what about criminal law?

- Criminal law: covers actions deemed to be sufficiently bad that they may take away your liberty
  - E.g., Fraud, CFAA (hacking), Wiretap, Extortion, Identity Theft, possession of CSAM, etc.
- We shouldn’t do measurement studies that break criminal laws? Agreed?
- Whose laws?
  - Laws in country of researcher? In any country where measurement takes place? Does nationality of researcher matter?
  - When the West does censorship evasion research are they careful to check that they aren’t breaking any laws in Iran, Russia or China (for example?)
- Ok, what if we just say that papers published in the US should not do things that would violate US criminal law?

Published at FOCI (co-located with PETS) in Washington D.C.

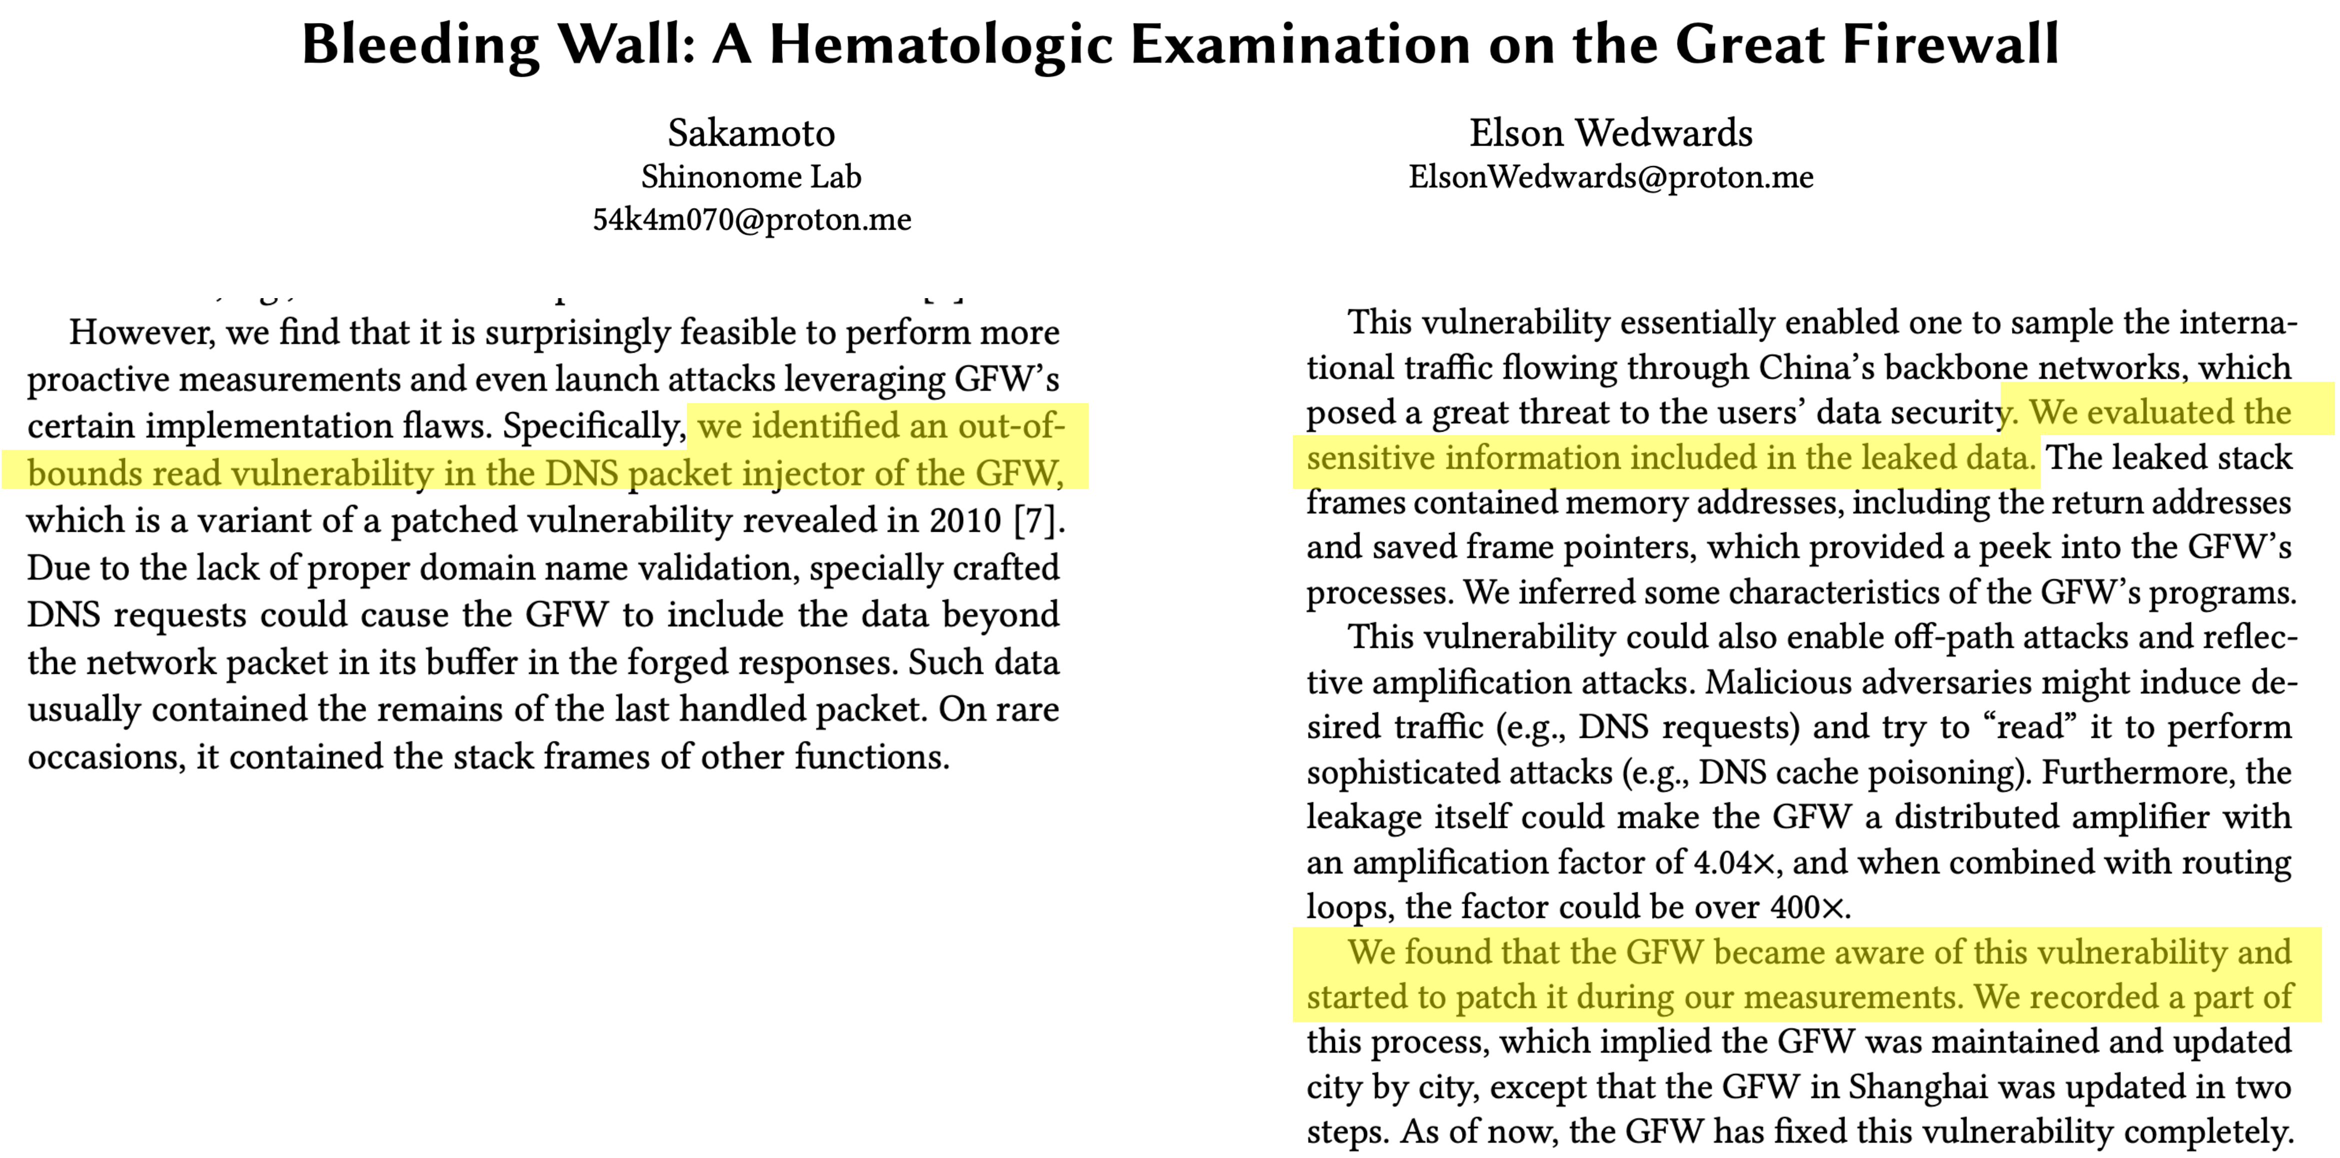

Bleeding Wall: A Hematologic Examination on the Great Firewall

Sakamoto
Shinonome Lab
54k4m070@proton.me

Elson Wedwards
ElsonWedwards@proton.me

However, we find that it is surprisingly feasible to perform more proactive measurements and even launch attacks leveraging GFW's certain implementation flaws. Specifically, we identified an out-of-bounds read vulnerability in the DNS packet injector of the GFW, which is a variant of a patched vulnerability revealed in 2010 [7]. Due to the lack of proper domain name validation, specially crafted DNS requests could cause the GFW to include the data beyond the network packet in its buffer in the forged responses. Such data usually contained the remains of the last handled packet. On rare occasions, it contained the stack frames of other functions.

This vulnerability essentially enabled one to sample the international traffic flowing through China's backbone networks, which posed a great threat to the users' data security. We evaluated the sensitive information included in the leaked data. The leaked stack frames contained memory addresses, including the return addresses and saved frame pointers, which provided a peek into the GFW's processes. We inferred some characteristics of the GFW's programs.

This vulnerability could also enable off-path attacks and reflective amplification attacks. Malicious adversaries might induce desired traffic (e.g., DNS requests) and try to "read" it to perform sophisticated attacks (e.g., DNS cache poisoning). Furthermore, the leakage itself could make the GFW a distributed amplifier with an amplification factor of 4.04×, and when combined with routing loops, the factor could be over 400×.

We found that the GFW became aware of this vulnerability and started to patch it during our measurements. We recorded a part of this process, which implied the GFW was maintained and updated city by city, except that the GFW in Shanghai was updated in two steps. As of now, the GFW has fixed this vulnerability completely.

Can We Sniff Wi-Fi?

Implications of *Joffe v. Google*

Simson L. Garfinkel and Michael McCarrin | Naval Postgraduate School

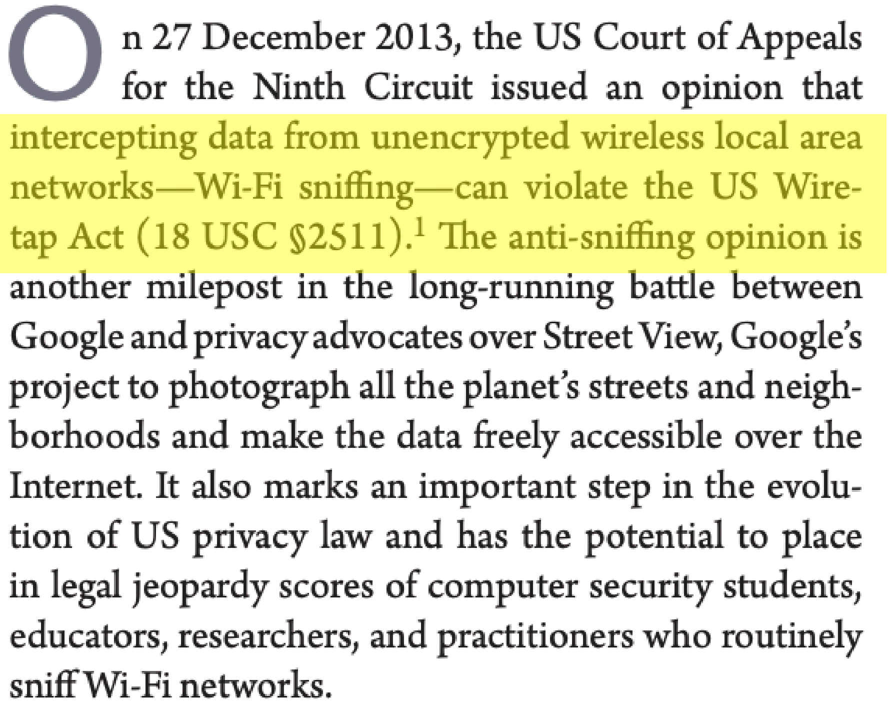

On 27 December 2013, the US Court of Appeals for the Ninth Circuit issued an opinion that intercepting data from unencrypted wireless local area networks—Wi-Fi sniffing—can violate the US Wiretap Act (18 USC §2511).[1] The anti-sniffing opinion is another milepost in the long-running battle between Google and privacy advocates over Street View, Google's project to photograph all the planet's streets and neighborhoods and make the data freely accessible over the Internet. It also marks an important step in the evolution of US privacy law and has the potential to place in legal jeopardy scores of computer security students, educators, researchers, and practitioners who routinely sniff Wi-Fi networks.

- This is true and it is "good law" in the 9th circuit (which includes CA, WA)
- Do you think researchers in CA stopped sniffing WiFi? Or passive wireless monitoring of open networks in general?

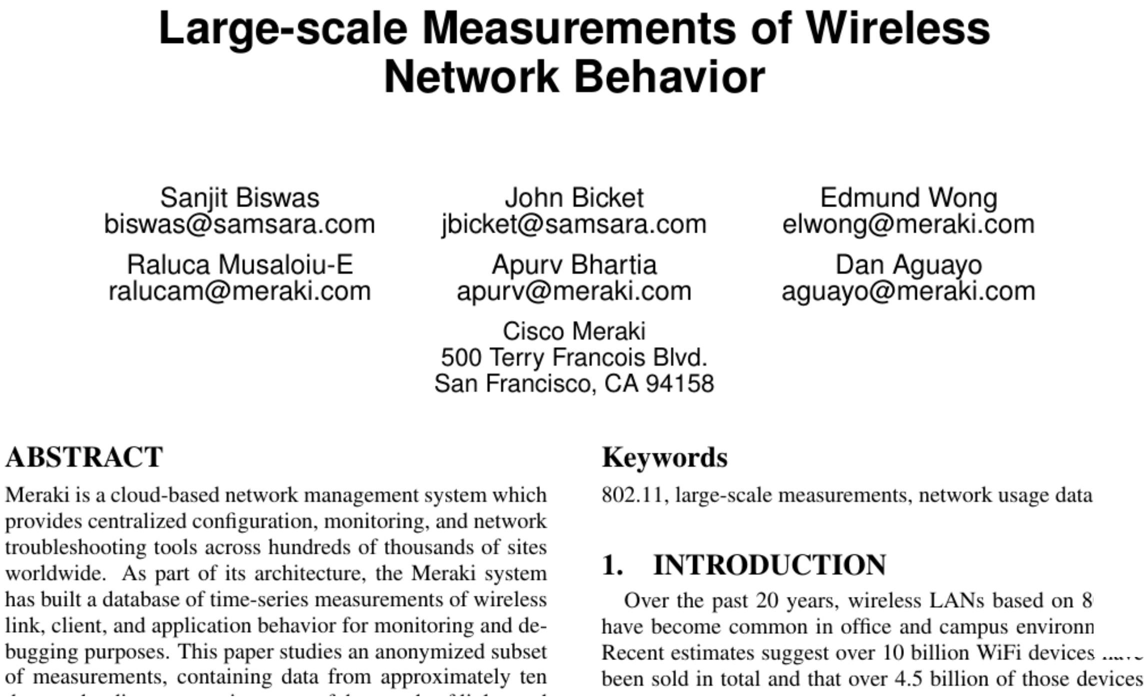

Large-scale Measurements of Wireless Network Behavior

Sanjit Biswas
biswas@samsara.com

John Bicket
jbicket@samsara.com

Edmund Wong
elwong@meraki.com

Raluca Musaloiu-E
ralucam@meraki.com

Apurv Bhartia
apurv@meraki.com

Dan Aguayo
aguayo@meraki.com

Cisco Meraki
500 Terry Francois Blvd.
San Francisco, CA 94158

ABSTRACT

Meraki is a cloud-based network management system which provides centralized configuration, monitoring, and network troubleshooting tools across hundreds of thousands of sites worldwide. As part of its architecture, the Meraki system has built a database of time-series measurements of wireless link, client, and application behavior for monitoring and debugging purposes. This paper studies an anonymized subset of measurements, containing data from approximately ten

Keywords

802.11, large-scale measurements, network usage data

1. INTRODUCTION

Over the past 20 years, wireless LANs based on 8
have become common in office and campus environn
Recent estimates suggest over 10 billion WiFi devices have
been sold in total and that over 4.5 billion of those devices

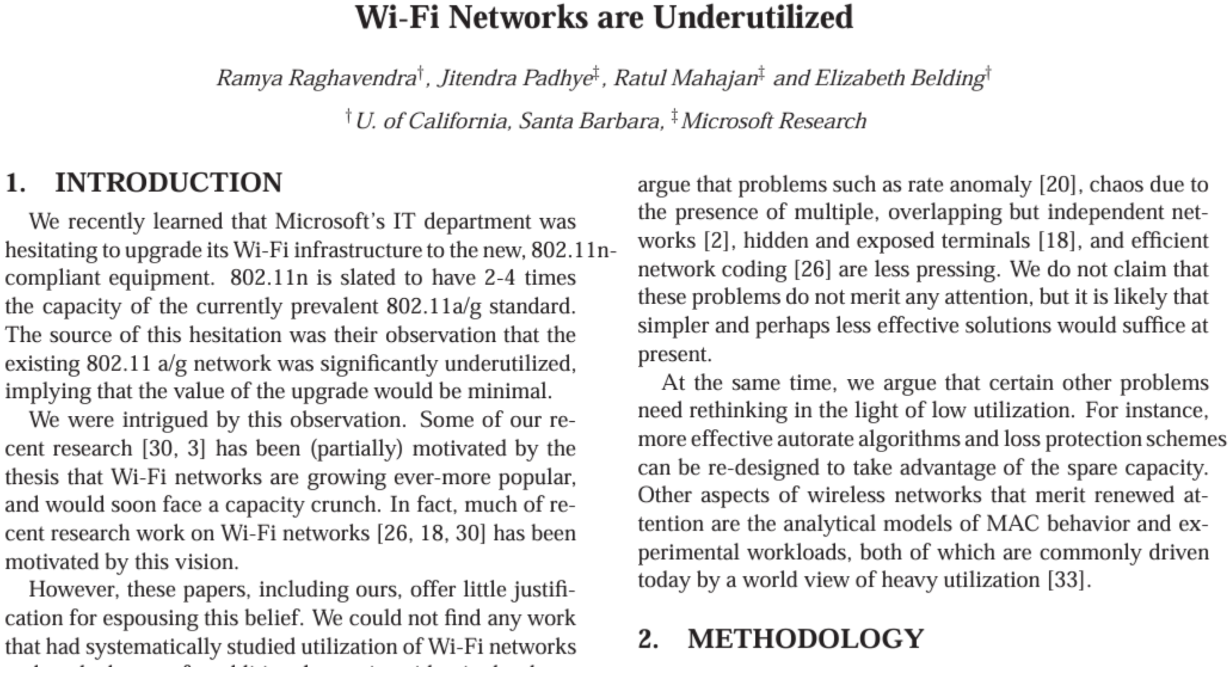

Wi-Fi Networks are Underutilized

Ramya Raghavendra[†], *Jitendra Padhye*[‡], *Ratul Mahajan*[‡] *and Elizabeth Belding*[†]

[†]*U. of California, Santa Barbara,* [‡]*Microsoft Research*

1. INTRODUCTION

We recently learned that Microsoft's IT department was hesitating to upgrade its Wi-Fi infrastructure to the new, 802.11n-compliant equipment. 802.11n is slated to have 2-4 times the capacity of the currently prevalent 802.11a/g standard. The source of this hesitation was their observation that the existing 802.11 a/g network was significantly underutilized, implying that the value of the upgrade would be minimal.

We were intrigued by this observation. Some of our recent research [30, 3] has been (partially) motivated by the thesis that Wi-Fi networks are growing ever-more popular, and would soon face a capacity crunch. In fact, much of recent research work on Wi-Fi networks [26, 18, 30] has been motivated by this vision.

However, these papers, including ours, offer little justification for espousing this belief. We could not find any work that had systematically studied utilization of Wi-Fi networks

argue that problems such as rate anomaly [20], chaos due to the presence of multiple, overlapping but independent networks [2], hidden and exposed terminals [18], and efficient network coding [26] are less pressing. We do not claim that these problems do not merit any attention, but it is likely that simpler and perhaps less effective solutions would suffice at present.

At the same time, we argue that certain other problems need rethinking in the light of low utilization. For instance, more effective autorate algorithms and loss protection schemes can be re-designed to take advantage of the spare capacity. Other aspects of wireless networks that merit renewed attention are the analytical models of MAC behavior and experimental workloads, both of which are commonly driven today by a world view of heavy utilization [33].

2. METHODOLOGY

- Why no change in researcher methodology or publication issues?

Where does that leave us wrt legality?

- Actual **cases** against researchers are incredibly rare (almost all civil)
- There are some baseline norms here wrt PCs
  - A few things are just out: no CSAM (even computer generated)
  - In general, actions that truly could harm a person or service (e.g., DDoS) get viewed through the legal lens (but some grey areas – e.g., Ben Zhou paper on traffic updates)
  - Completely unauthorized access (e.g., using stolen password, vulnerability, etc.) is usually seen through legal lens
    - Some exceptions: e.g., GFW paper from before
    - What about access to criminal infrastructure? E.g. spamalytics paper?
      How would that paper have been seen if we'd infiltrated a commercial p2p network
      and took over 2.5% of its traffic?
    - Is there a double standard against criminals?
  - Practical advice: contact general counsel before doing anything edgy
- BTW, will be offering the Cybersec and US Law 291 next quarter

Human Subjects:
Job of the Institutional Review Board (IRB)

- Minimize unnecessary risks to human subjects
  - Approve, require modifications in, or disapprove research involving human subjects *before* experimentation takes place
  - Safeguards for privacy and for vulnerable populations
  - Informed consent (when required)
  - Balance potential harms with likely benefits
  - Cynical version: ensure institutional compliance to avoid regulatory liability
- Particularly key in biomedical research
  - Some institutions have separate IRBs for biomed vs other research, but many have just one
- Institutions can have additional policies, but HHS regulations provide a baseline decision framework

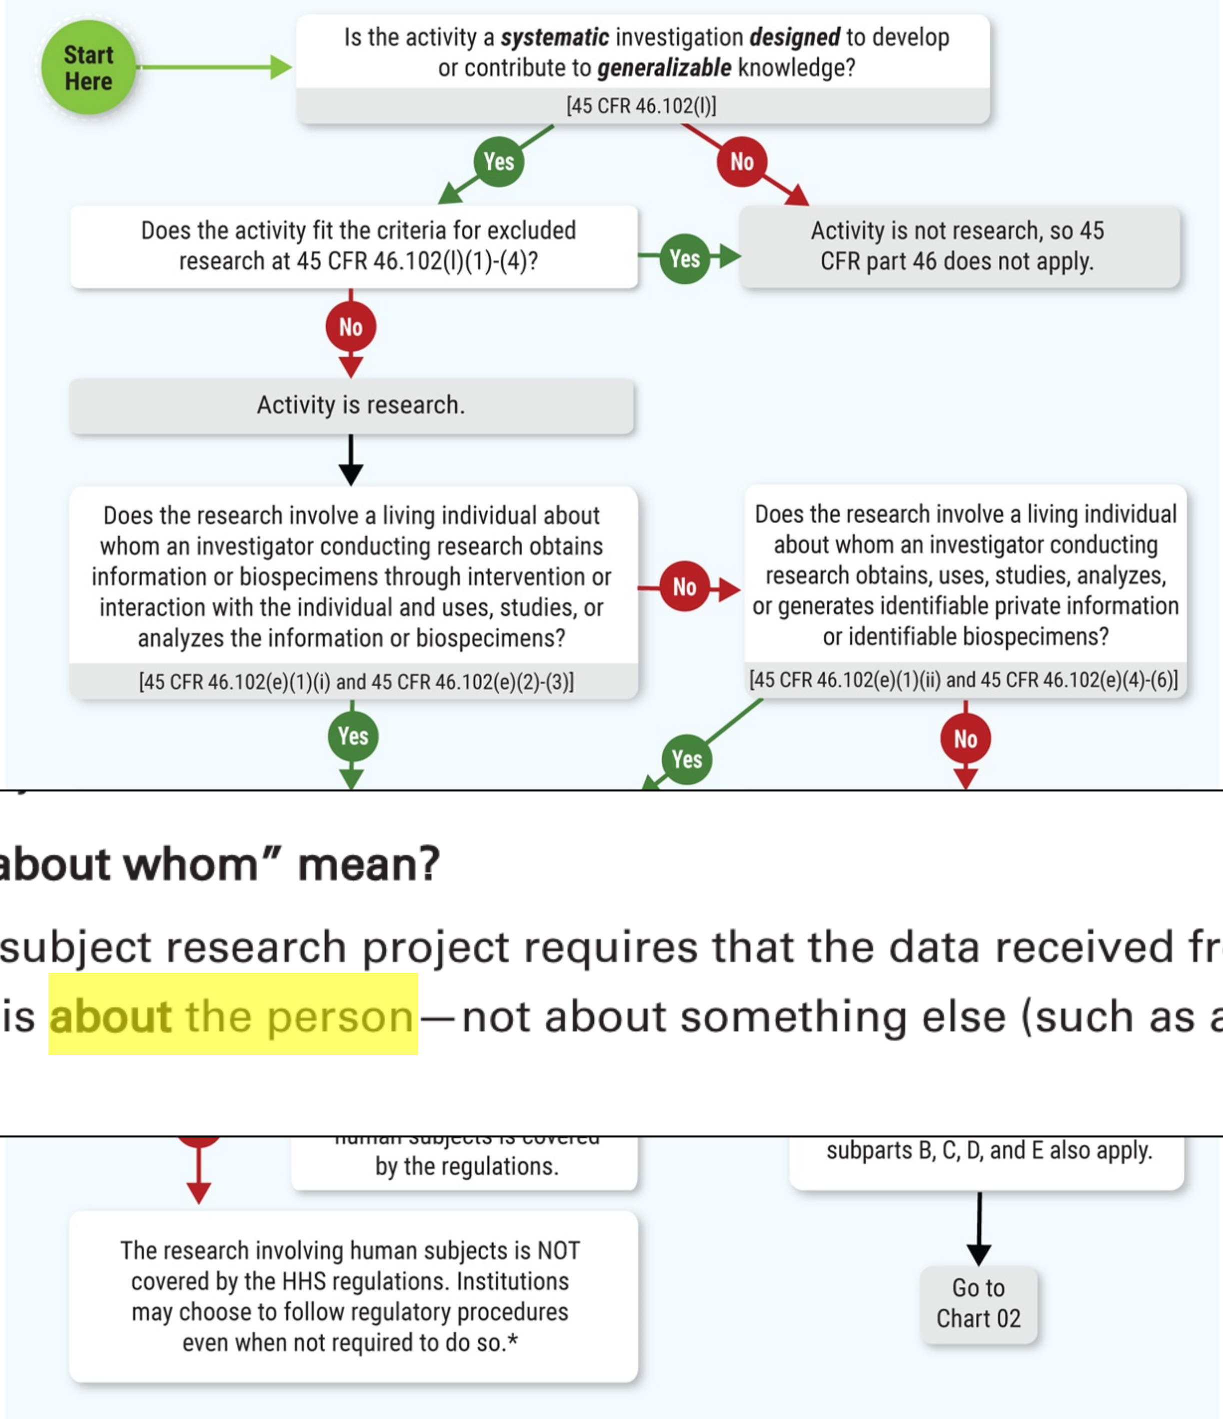

Start Here

Is the activity a ***systematic*** investigation ***designed*** to develop or contribute to ***generalizable*** knowledge?

[45 CFR 46.102(l)]

Yes

No

Does the activity fit the criteria for excluded research at 45 CFR 46.102(l)(1)-(4)?

Yes

Activity is not research, so 45 CFR part 46 does not apply.

No

Activity is research.

Does the research involve a living individual about whom an investigator conducting research obtains information or biospecimens through intervention or interaction with the individual and uses, studies, or analyzes the information or biospecimens?

[45 CFR 46.102(e)(1)(i) and 45 CFR 46.102(e)(2)-(3)]

No

Does the research involve a living individual about whom an investigator conducting research obtains, uses, studies, analyzes, or generates identifiable private information or identifiable biospecimens?

[45 CFR 46.102(e)(1)(ii) and 45 CFR 46.102(e)(4)-(6)]

Yes

Yes

No

human subjects is covered by the regulations.

subparts B, C, D, and E also apply.

The research involving human subjects is NOT covered by the HHS regulations. Institutions may choose to follow regulatory procedures even when not required to do so.*

Go to Chart 02

- **What does "about whom" mean?**
  - A human subject research project requires that the data received from the living individual is about the person—not about something else (such as a product or service).

Is it human subjects research?

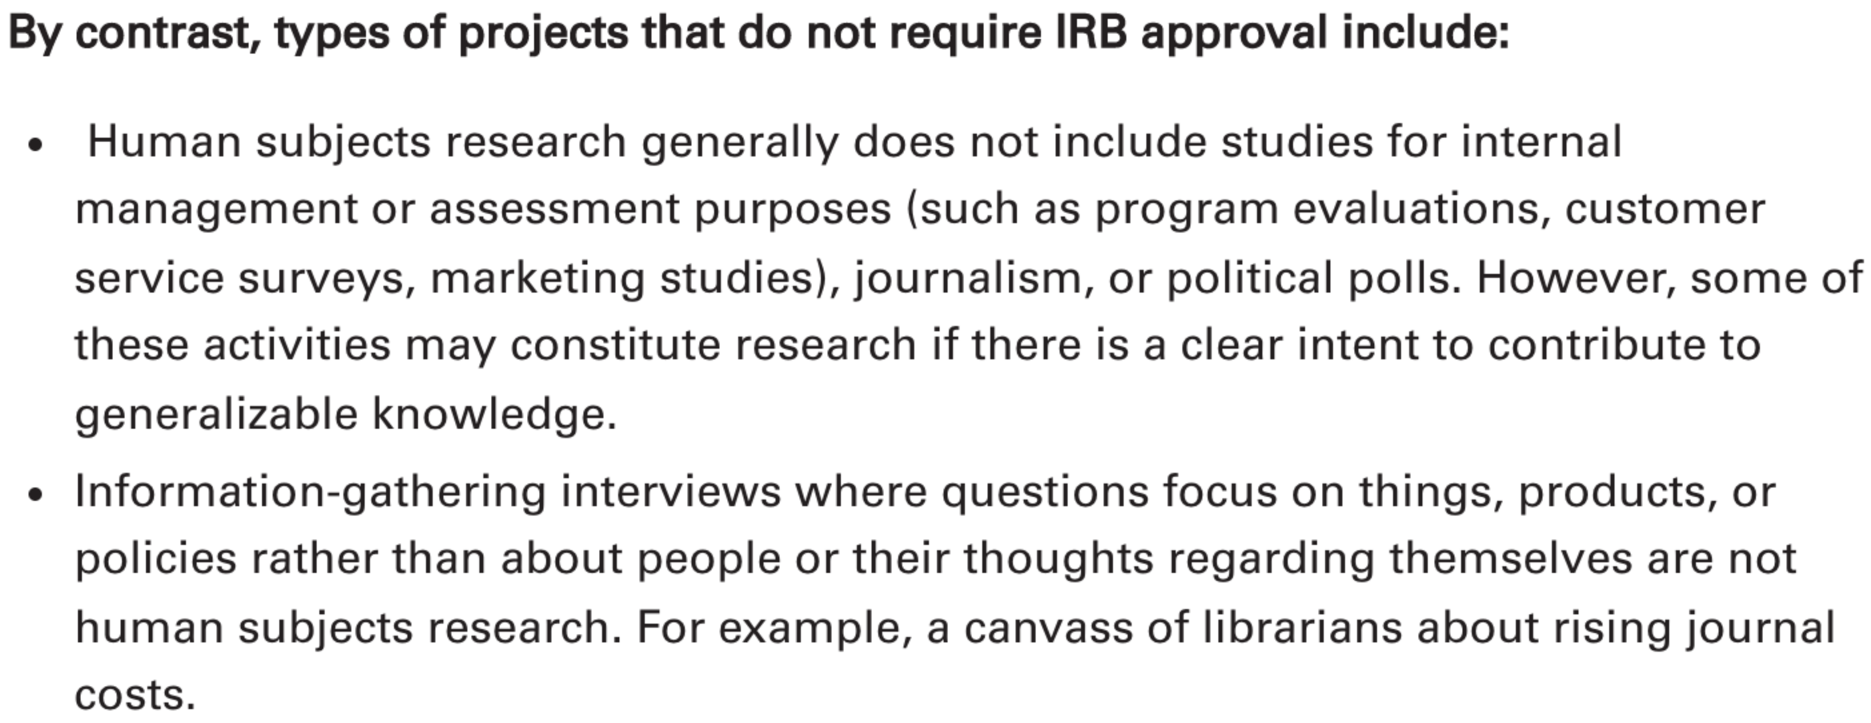

By contrast, types of projects that do not require IRB approval include:

- Human subjects research generally does not include studies for internal management or assessment purposes (such as program evaluations, customer service surveys, marketing studies), journalism, or political polls. However, some of these activities may constitute research if there is a clear intent to contribute to generalizable knowledge.
- Information-gathering interviews where questions focus on things, products, or policies rather than about people or their thoughts regarding themselves are not human subjects research. For example, a canvass of librarians about rising journal costs.

What are the implications of this definition?

- Which is human subjects research?
  - You install workplace monitors
    - to see when and where employees use space to guide space allocation and capital investment on new buildings?
    - to characterize the different modalities and patterns of use?
  - You acquire and published detailed personal information about Russian cybercriminals
    - As part of a paper at USENIX Security characterizing the tactics of Russian cybercriminals
    - As part of a newspaper feature series about the rise of Russian cybercriminal gangs
  - You do a study where you call tech company representatives, you explain (deceiving them) that you want to do a security analysis on one of their products and will they give permission?
  - You call Amazon tech support workers and ask them how they feel about their jobs?

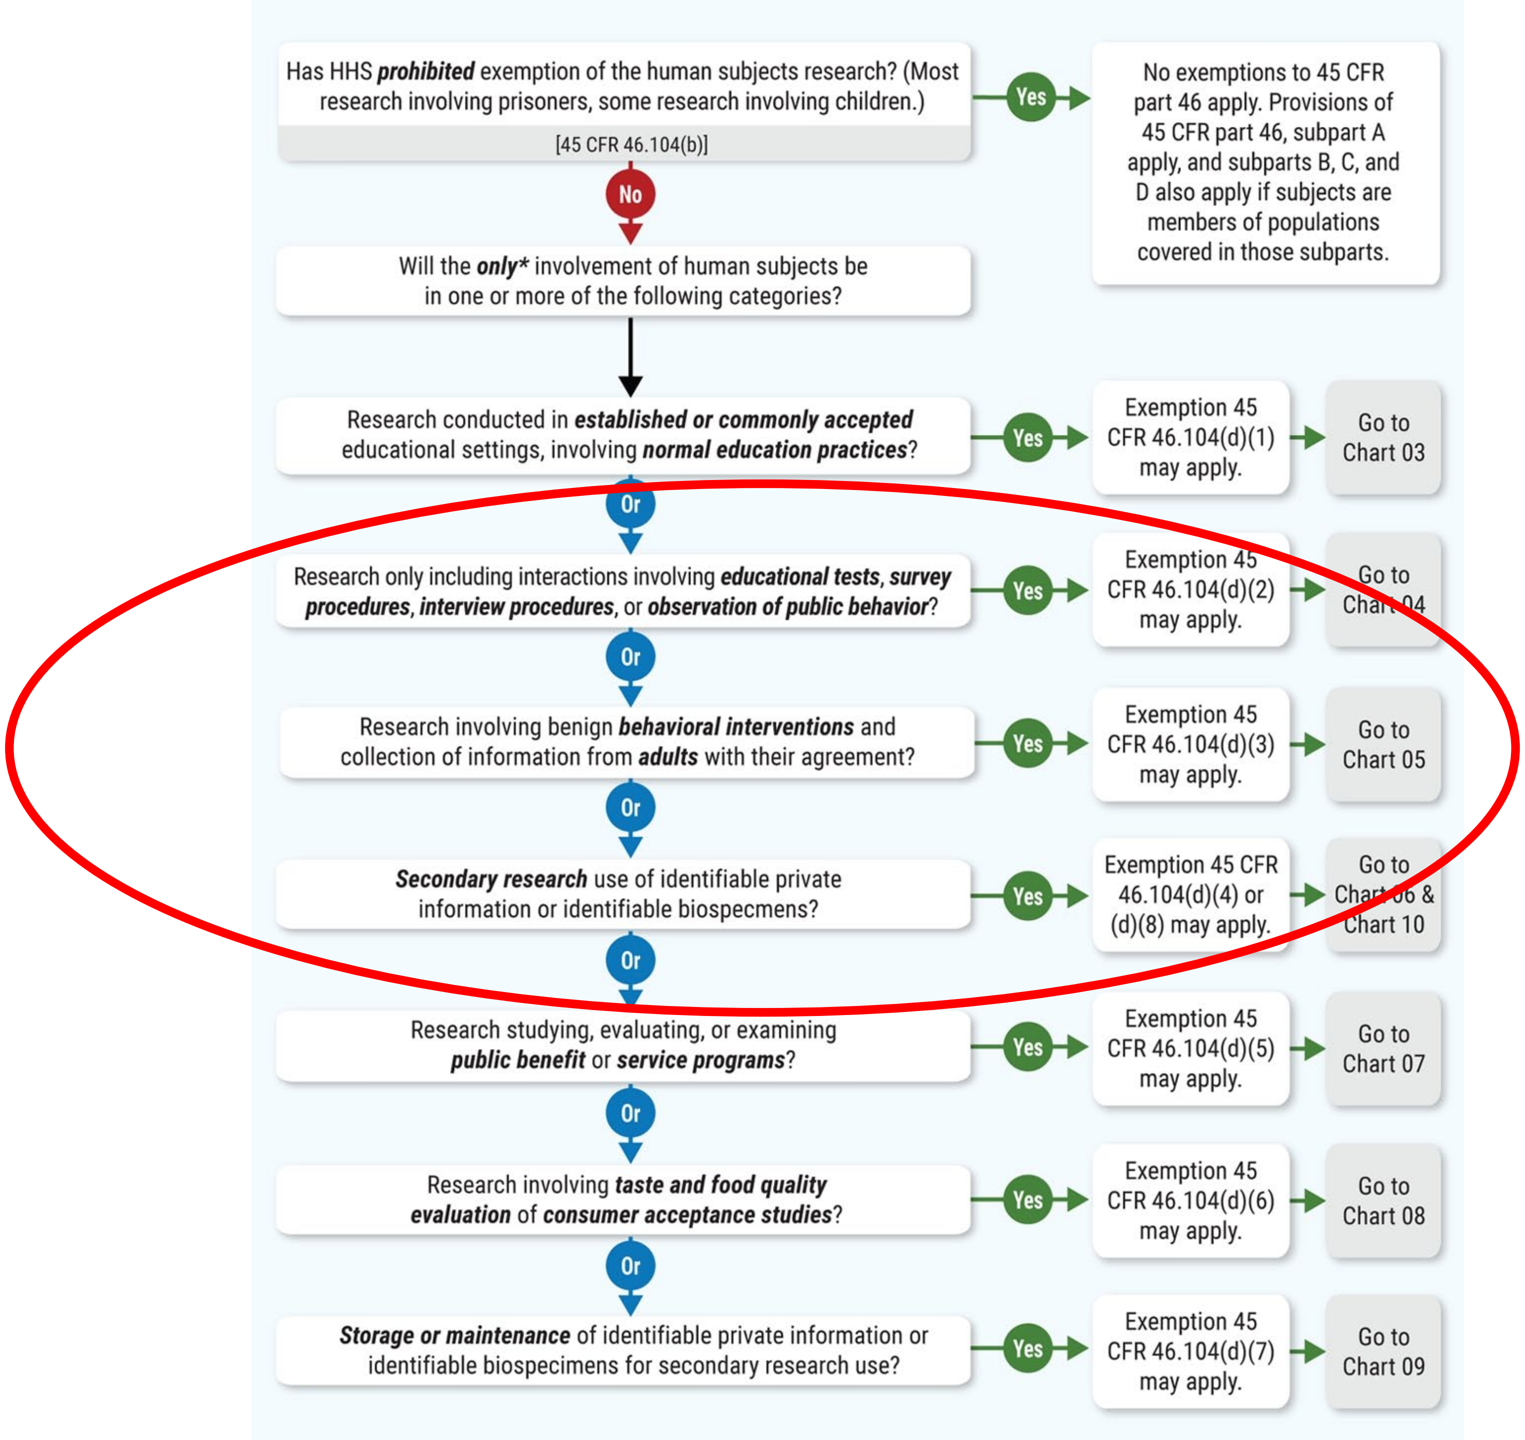

Has HHS ***prohibited*** exemption of the human subjects research? (Most research involving prisoners, some research involving children.)

[45 CFR 46.104(b)]

Yes → No exemptions to 45 CFR part 46 apply. Provisions of 45 CFR part 46, subpart A apply, and subparts B, C, and D also apply if subjects are members of populations covered in those subparts.

No ↓

Will the ***only**** involvement of human subjects be in one or more of the following categories?

Research conducted in ***established or commonly accepted*** educational settings, involving ***normal education practices***?

Yes → Exemption 45 CFR 46.104(d)(1) may apply. → Go to Chart 03

Or

Research only including interactions involving ***educational tests***, ***survey procedures***, ***interview procedures***, or ***observation of public behavior***?

Yes → Exemption 45 CFR 46.104(d)(2) may apply. → Go to Chart 04

Or

Research involving benign ***behavioral interventions*** and collection of information from ***adults*** with their agreement?

Yes → Exemption 45 CFR 46.104(d)(3) may apply. → Go to Chart 05

Or

Secondary research use of identifiable private information or identifiable biospecmens?

Yes → Exemption 45 CFR 46.104(d)(4) or (d)(8) may apply. → Go to Chart 06 & Chart 10

Or

Research studying, evaluating, or examining ***public benefit*** or ***service programs***?

Yes → Exemption 45 CFR 46.104(d)(5) may apply. → Go to Chart 07

Or

Research involving ***taste and food quality evaluation*** of ***consumer acceptance studies***?

Yes → Exemption 45 CFR 46.104(d)(6) may apply. → Go to Chart 08

Or

Storage or maintenance of identifiable private information or identifiable biospecimens for secondary research use?

Yes → Exemption 45 CFR 46.104(d)(7) may apply. → Go to Chart 09

General rules

- Not human subjects research – no limitations placed by IRB
- Exempt category – few limitations if compliant
- Otherwise
  - Document protocol and controls against injury to subjects and argue why benefits exceed risks
  - Prior consent
  - Post-experiment debrief
  - Note waivers possible for these last two
    - Typically need to show necessity or that harm would arise as a result
- IRB can request further controls/modifications

Quick practical discussion

- You **always** want to get IRB review if there is any chance that someone might question the ethics of your work
- But… being declared either “not human subjects” or “exempt” does not mean everyone will agree your work is ethical – just that it didn’t qualify for detailed human subjects controls

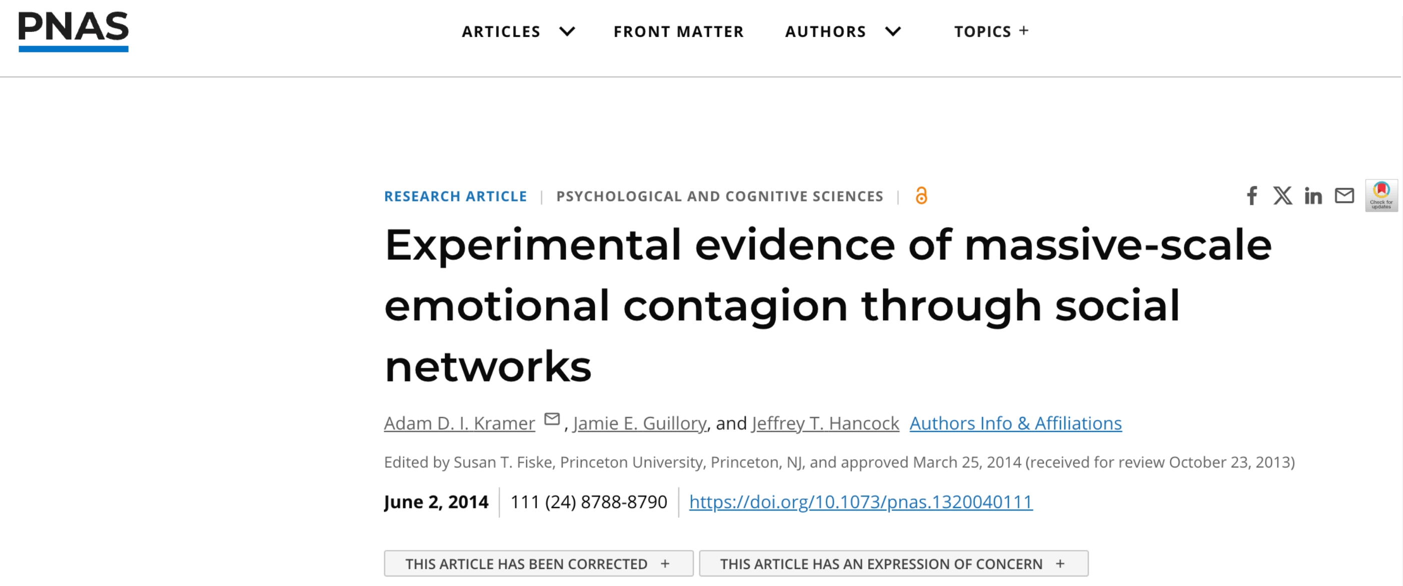

PNAS

ARTICLES FRONT MATTER AUTHORS TOPICS +

RESEARCH ARTICLE | PSYCHOLOGICAL AND COGNITIVE SCIENCES

Experimental evidence of massive-scale emotional contagion through social networks

Adam D. I. Kramer, Jamie E. Guillory, and Jeffrey T. Hancock Authors Info & Affiliations

Edited by Susan T. Fiske, Princeton University, Princeton, NJ, and approved March 25, 2014 (received for review October 23, 2013)

June 2, 2014 | 111 (24) 8788-8790 | https://doi.org/10.1073/pnas.1320040111

THIS ARTICLE HAS BEEN CORRECTED + | THIS ARTICLE HAS AN EXPRESSION OF CONCERN +

Facebook fiasco: was Cornell's study ‘emotional contagion’ an ethics brea

A covert experiment to influence the emotions of more than 600,000 people. A major scientific journal behaving like a rabbit in the headlights. A university in a PR tailspin

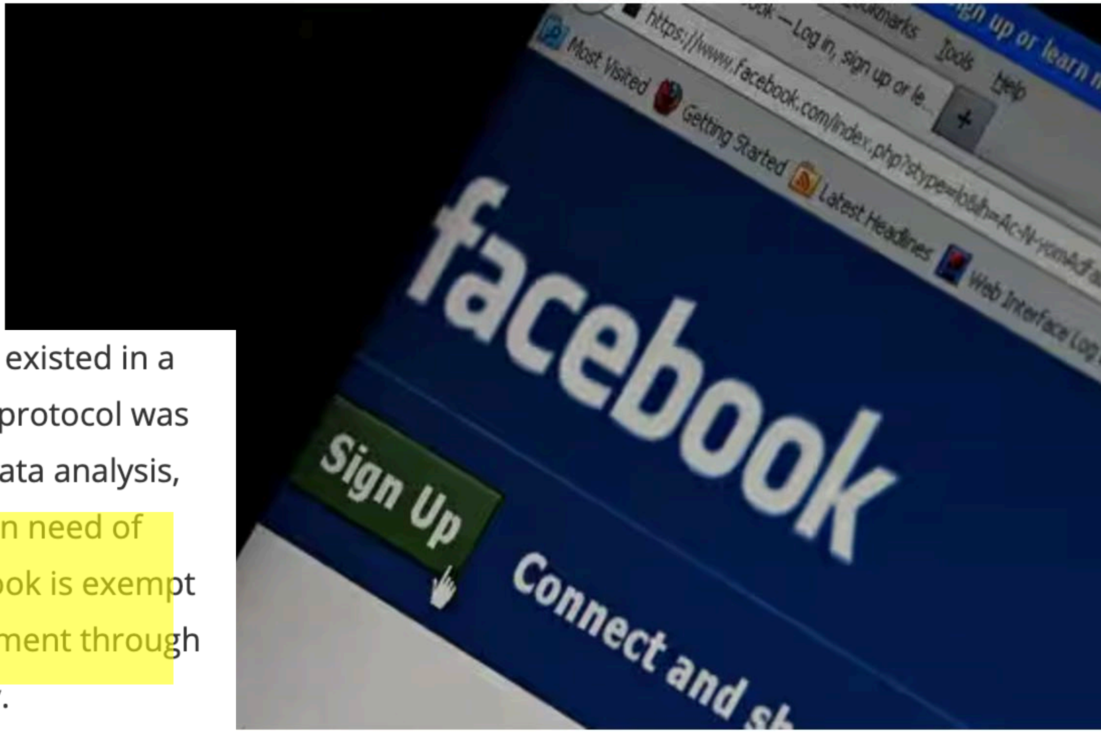

The experiment, as a collaboration between academic researchers at Cornell, and Facebook, existed in a grey area of federal regulation. It was allegedly designed by the Cornell researchers, but the protocol was carried out exclusively by Facebook employees. The Cornell researchers then assisted with data analysis, which the Cornell Institutional Review Board (IRB) deemed a “pre-existing data-set,” and not in need of substantive ethical review. As a private entity not in receipt of federal research funds, Facebook is exempt from review under the Common Rule. Though Facebook later claimed it reviewed the experiment through its own in-house oversight mechanisms, the exact structure of this review remains a mystery.

Facebook have recently come under fire for a controversial psychological study. Photog Dave Thompson/PA Photograph: Dave Thompson/PA

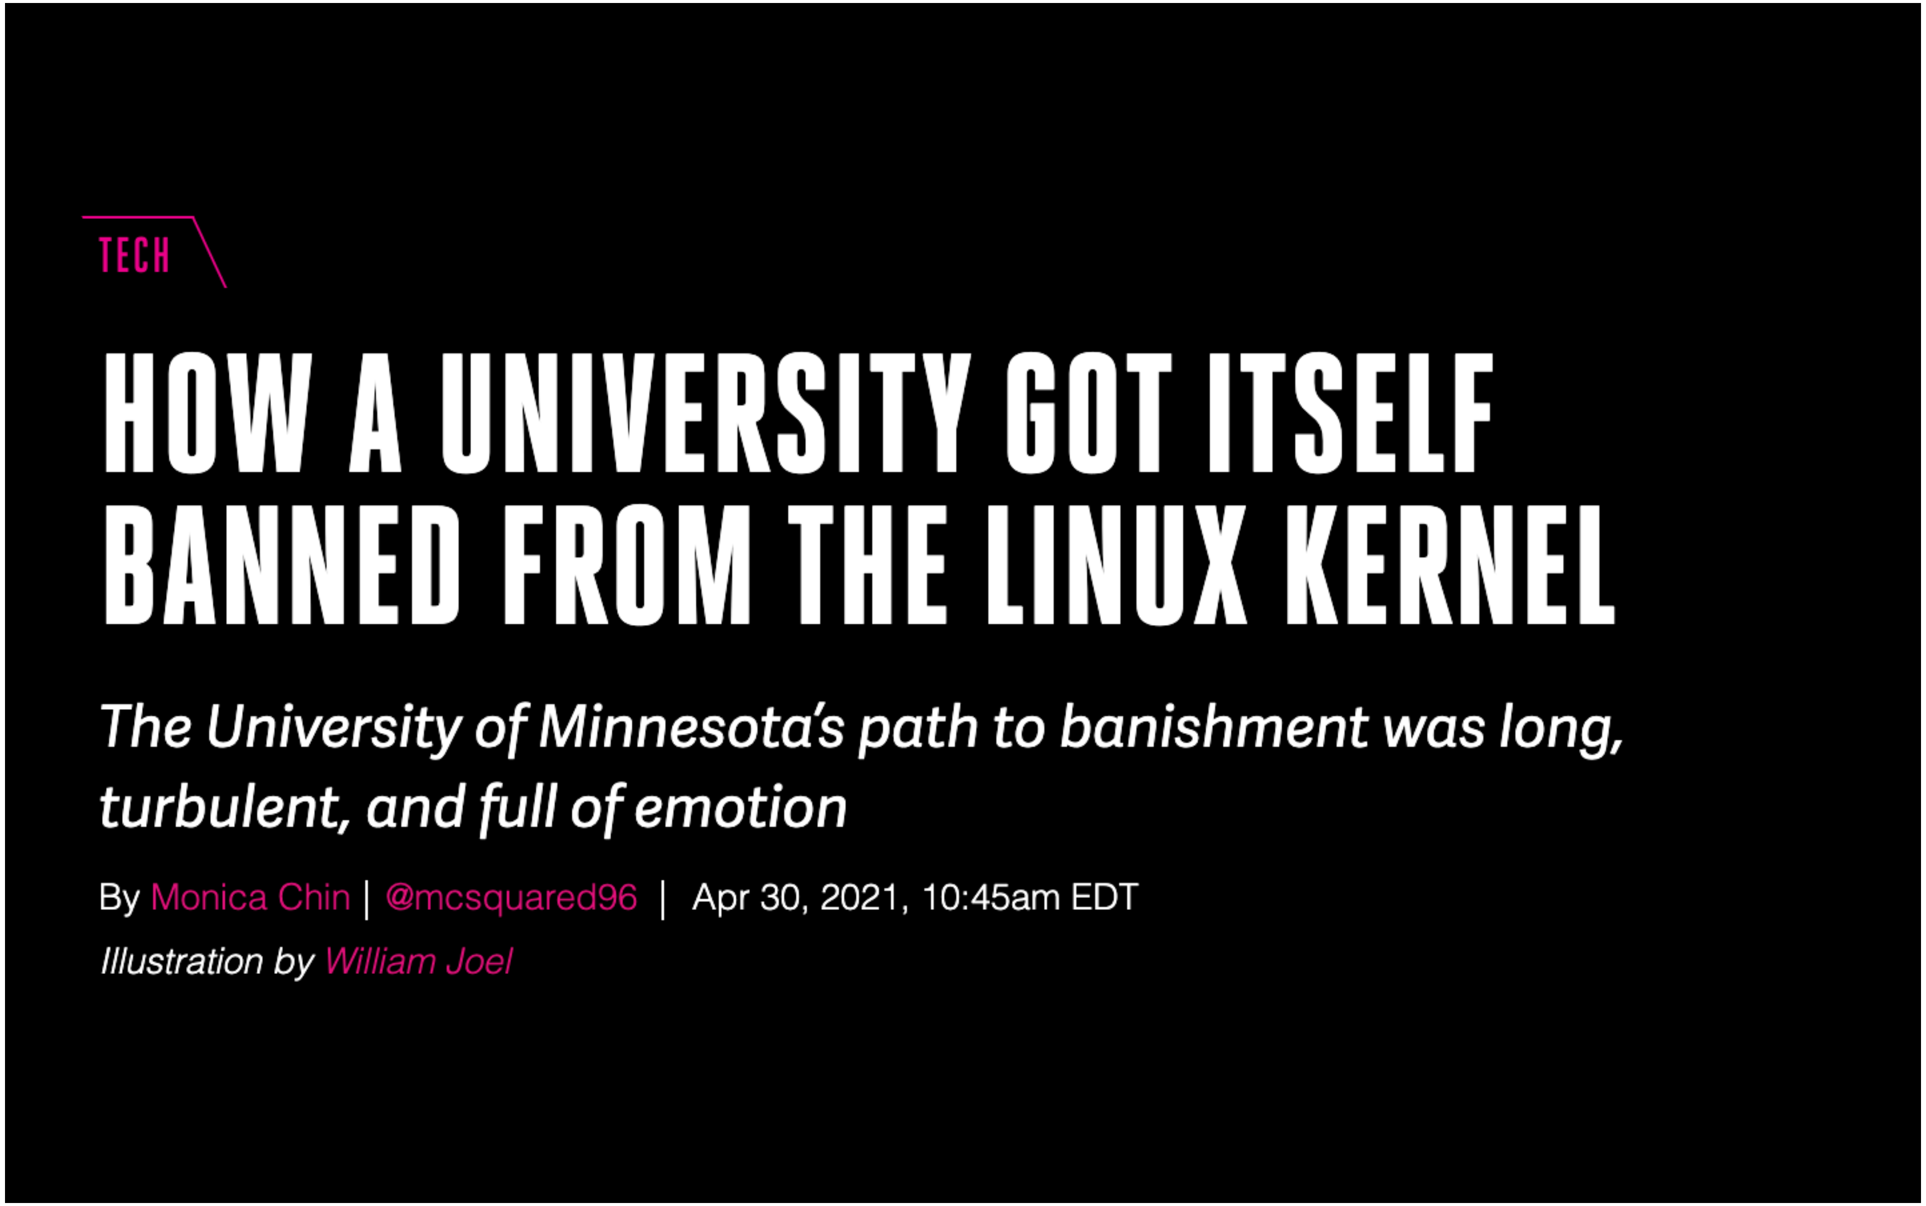

TECH

HOW A UNIVERSITY GOT ITSELF BANNED FROM THE LINUX KERNEL

The University of Minnesota's path to banishment was long, turbulent, and full of emotion

By Monica Chin | @mcsquared96 | Apr 30, 2021, 10:45am EDT

Illustration by William Joel

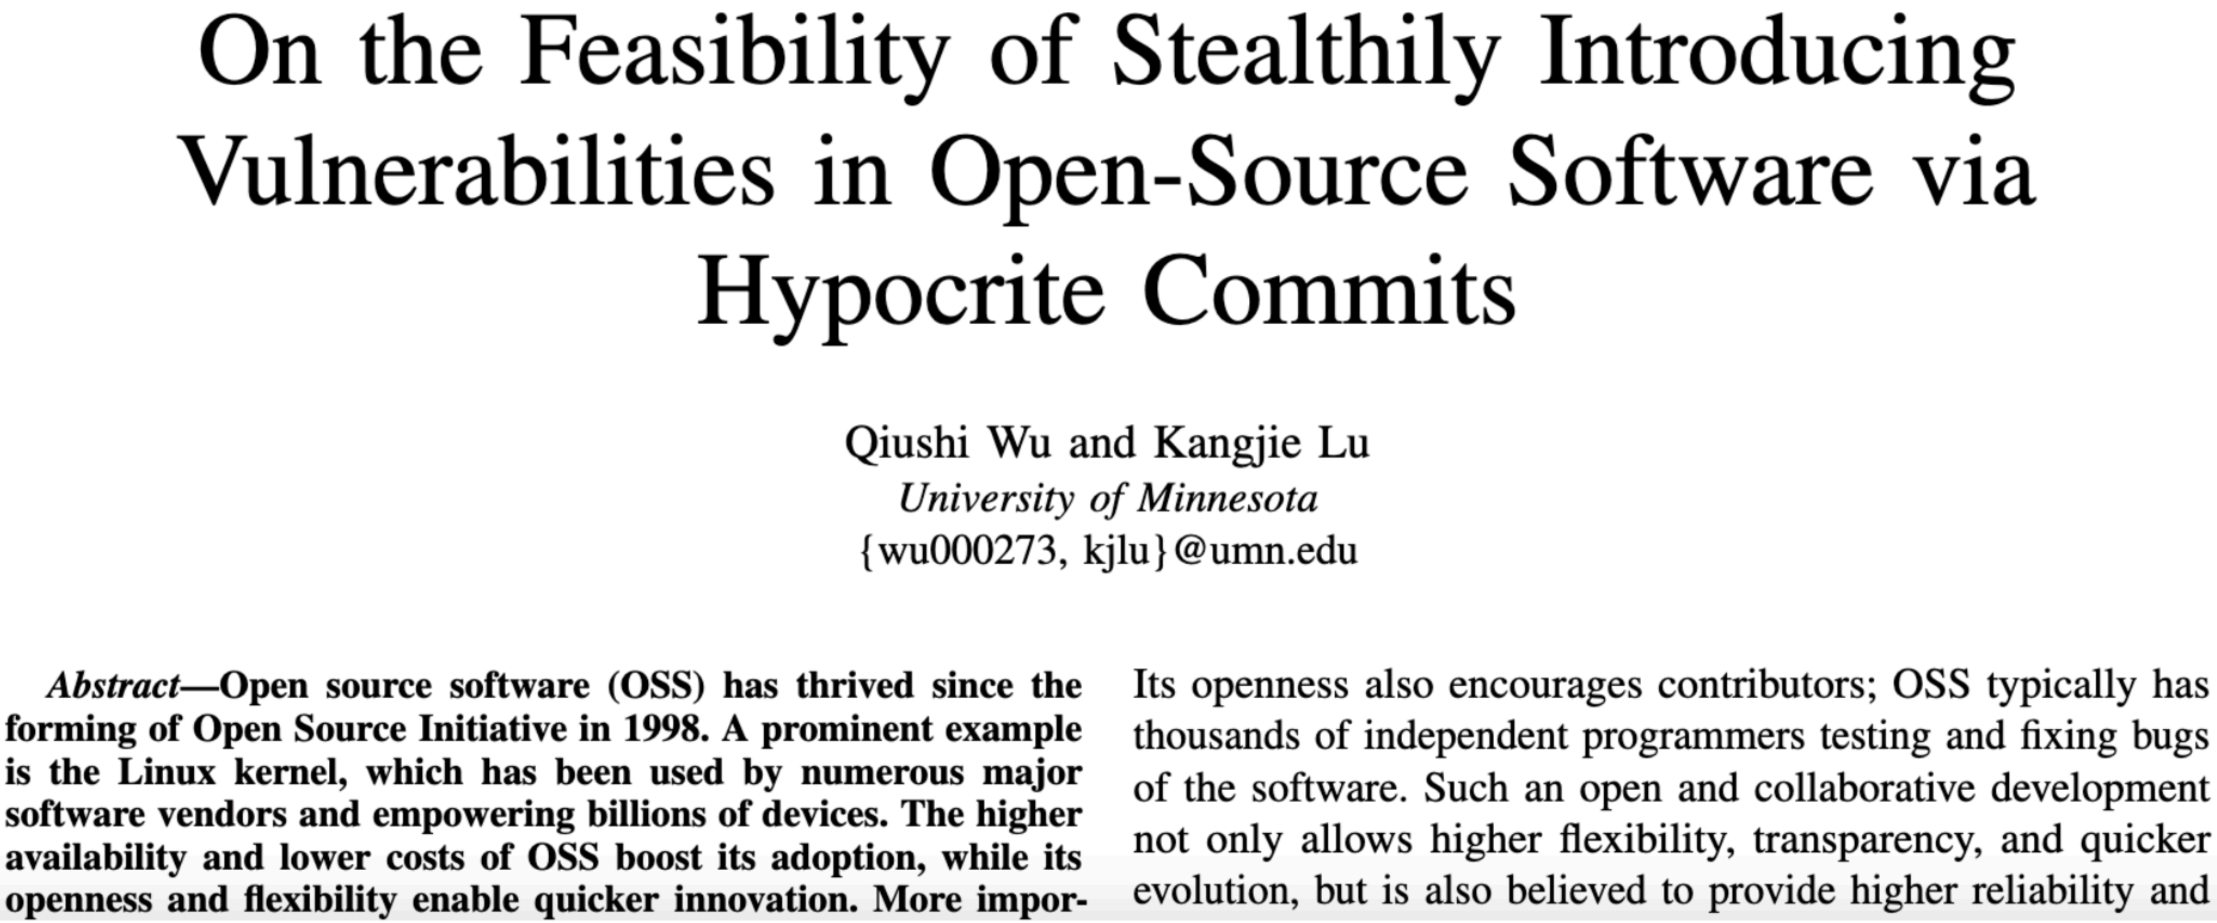

On the Feasibility of Stealthily Introducing Vulnerabilities in Open-Source Software via Hypocrite Commits

Qiushi Wu and Kangjie Lu
University of Minnesota
{wu000273, kjlu}@umn.edu

***Abstract*—Open source software (OSS) has thrived since the forming of Open Source Initiative in 1998. A prominent example is the Linux kernel, which has been used by numerous major software vendors and empowering billions of devices. The higher availability and lower costs of OSS boost its adoption, while its openness and flexibility enable quicker innovation. More impor-**

Its openness also encourages contributors; OSS typically has thousands of independent programmers testing and fixing bugs of the software. Such an open and collaborative development not only allows higher flexibility, transparency, and quicker evolution, but is also believed to provide higher reliability and

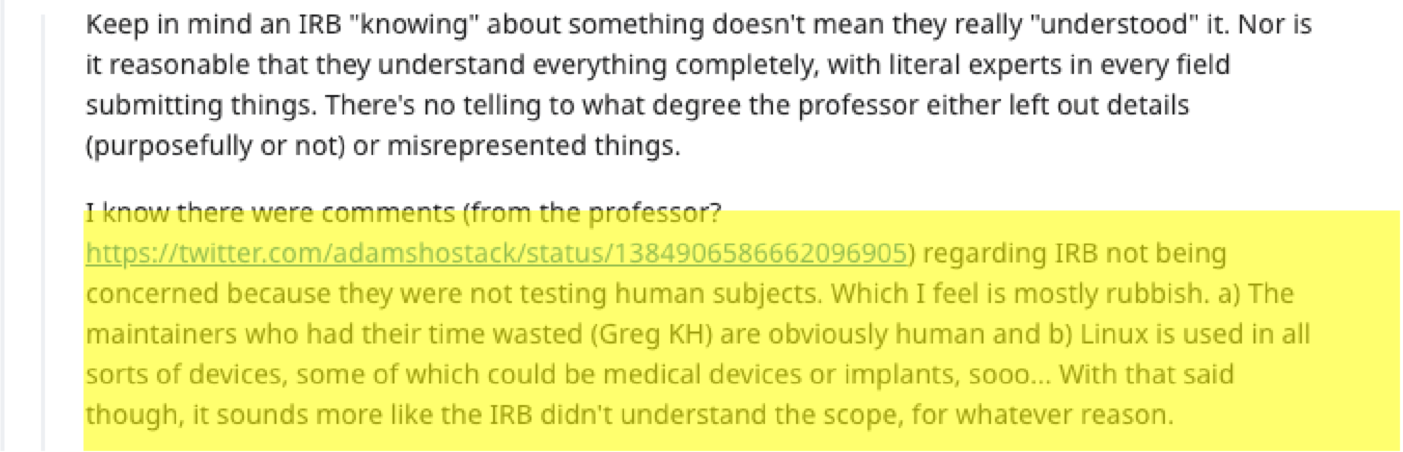

> Keep in mind an IRB "knowing" about something doesn't mean they really "understood" it. Nor is it reasonable that they understand everything completely, with literal experts in every field submitting things. There's no telling to what degree the professor either left out details (purposefully or not) or misrepresented things.
>
> I know there were comments (from the professor? https://twitter.com/adamshostack/status/1384906586662096905) regarding IRB not being concerned because they were not testing human subjects. Which I feel is mostly rubbish. a) The maintainers who had their time wasted (Greg KH) are obviously human and b) Linux is used in all sorts of devices, some of which could be medical devices or implants, sooo... With that said though, it sounds more like the IRB didn't understand the scope, for whatever reason.

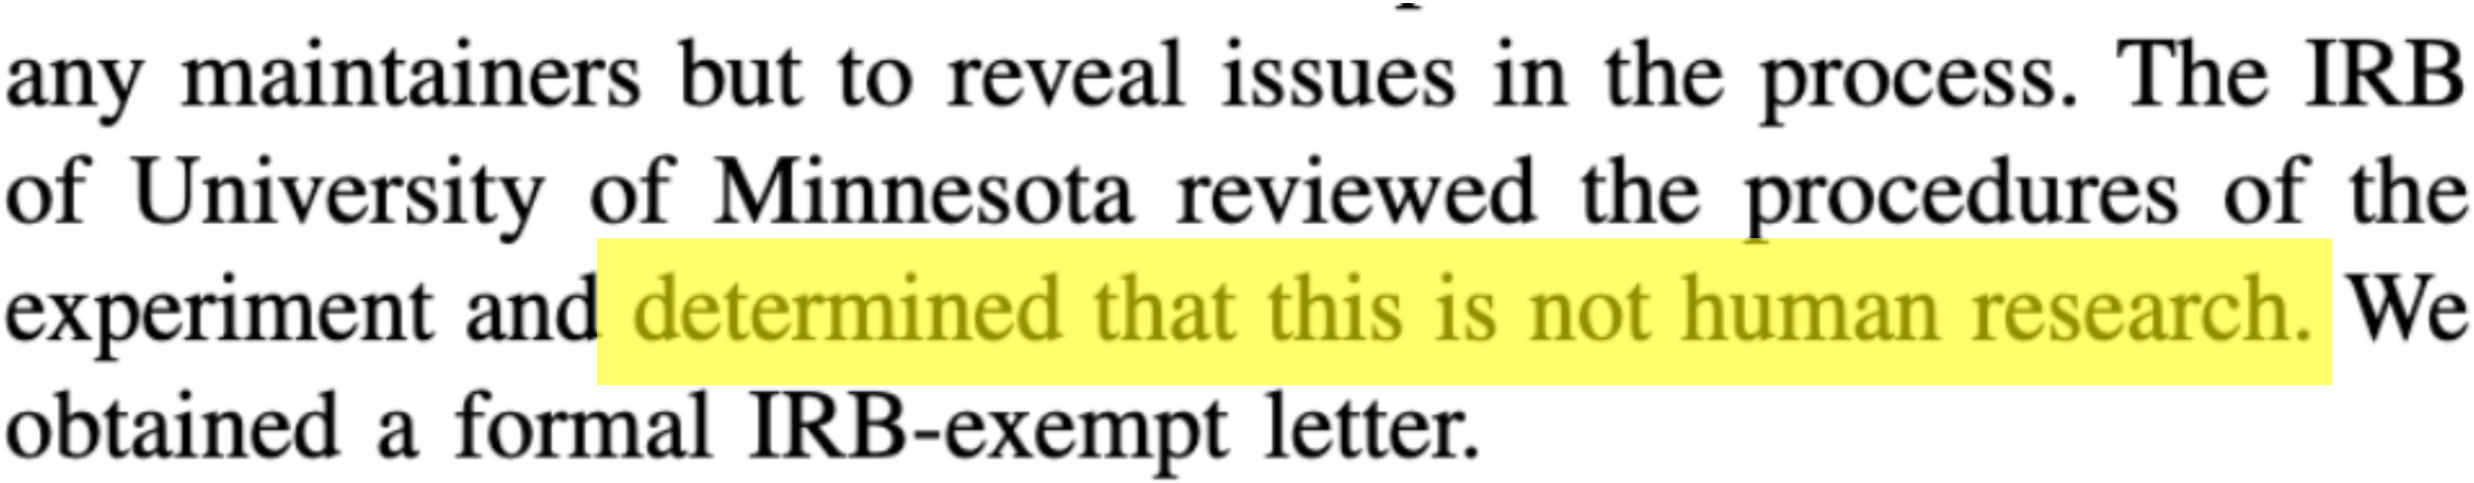

any maintainers but to reveal issues in the process. The IRB of University of Minnesota reviewed the procedures of the experiment and determined that this is not human research. We obtained a formal IRB-exempt letter.

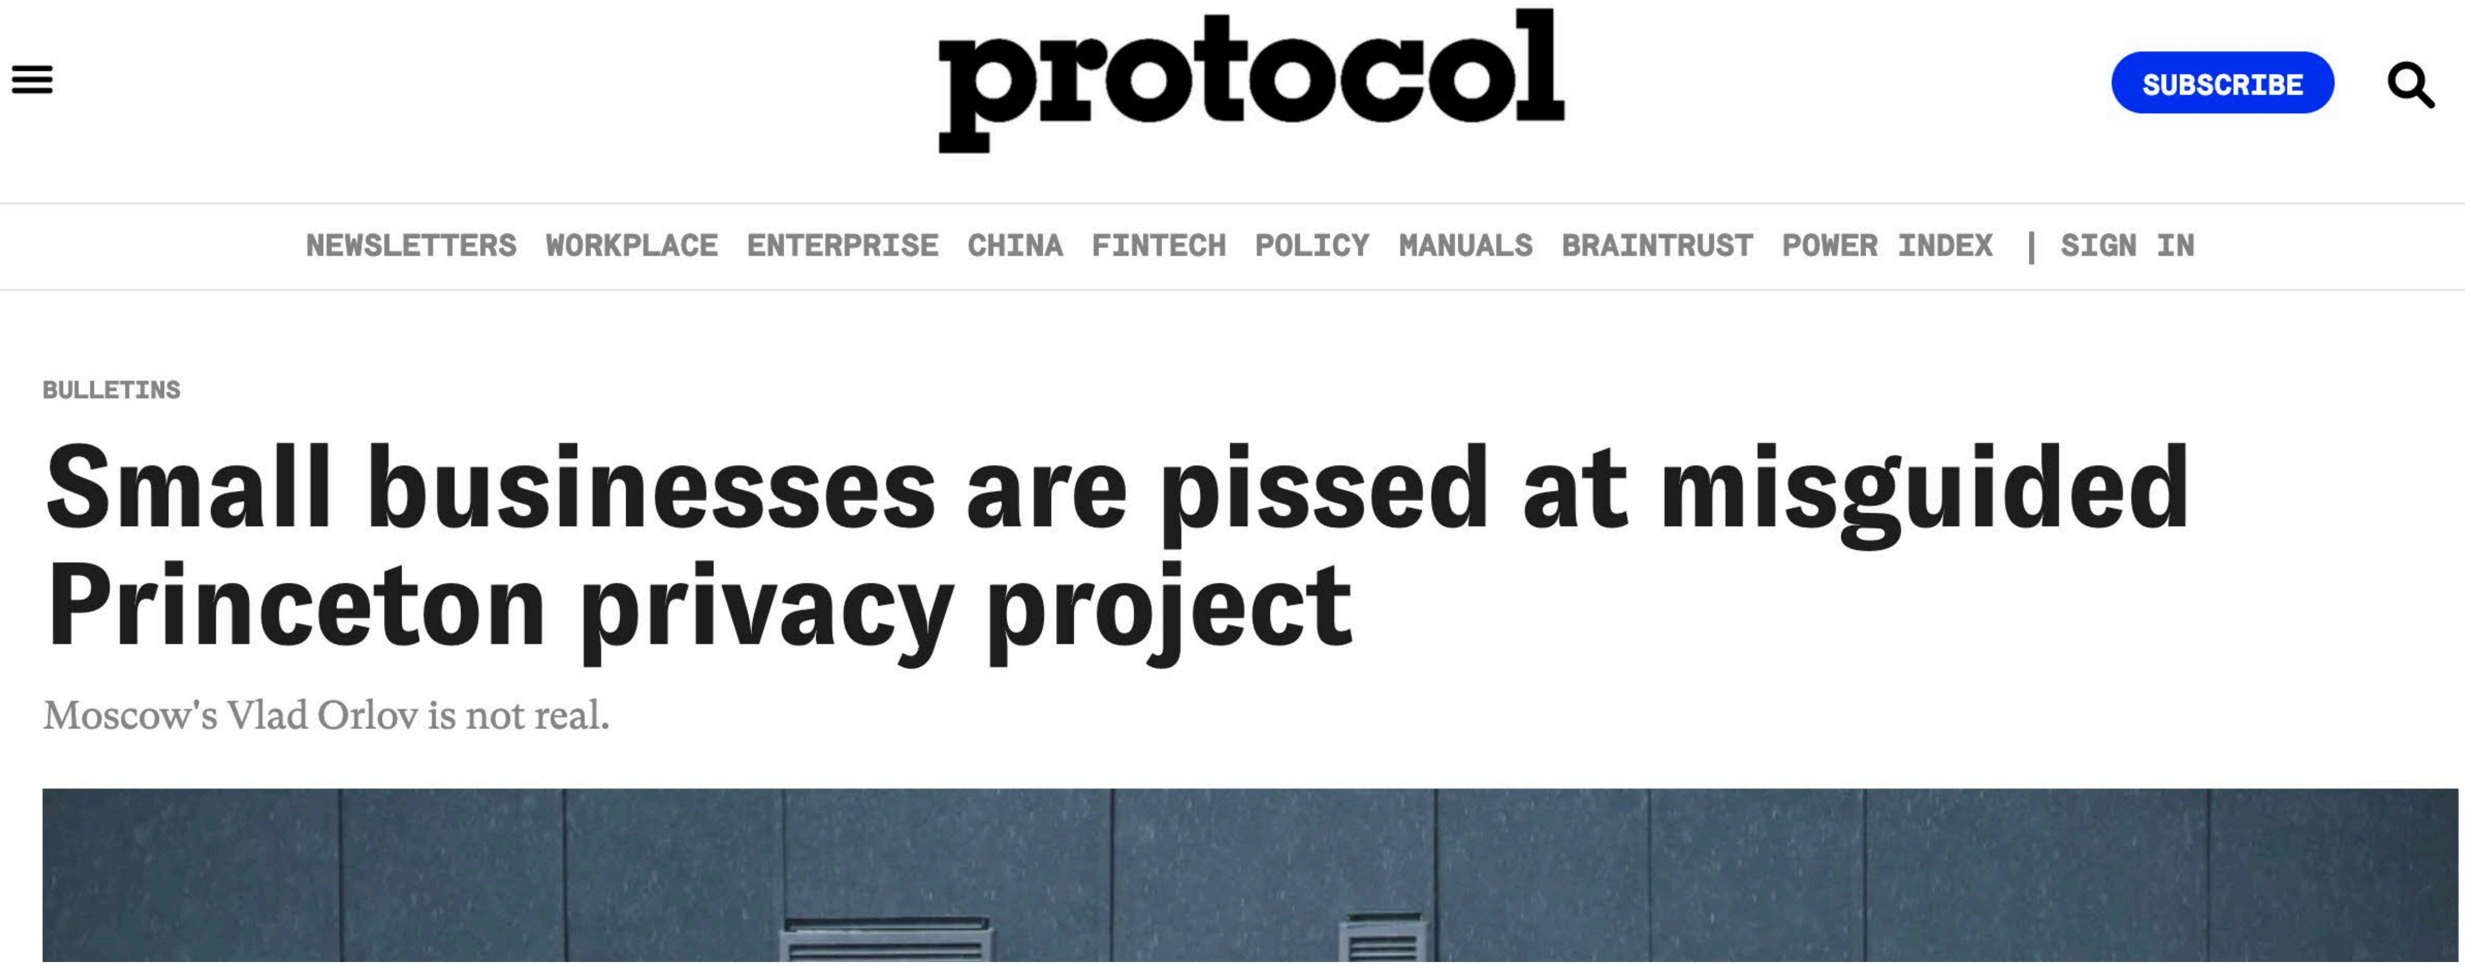

protocol
SUBSCRIBE
NEWSLETTERS WORKPLACE ENTERPRISE CHINA FINTECH POLICY MANUALS BRAINTRUST POWER INDEX | SIGN IN
BULLETINS

Small businesses are pissed at misguided Princeton privacy project

Moscow's Vlad Orlov is not real.

- **What are the goals of this research study?**

  The study aims to advance understanding of how websites have implemented the data rights provisions of European Union and California privacy law, specifically the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA).

  Our goals are to accurately describe how websites have operationalized these new user rights, whether websites are extending these rights to non-EU citizens and non-California residents, and whether websites are effectively authenticating users when they exercise these rights.

- **Why does this study involve contacting websites?**

  Very few websites post details of their processes for handling GDPR and CCPA requests. Both the GDPR and the CCPA contemplate users and intermediaries reaching out with questions about data rights processes, and we are using that opportunity to understand current website policies and practices.

- **Did an Institutional Review Board consider this study?**

  We submitted an application detailing our research methods to the Princeton University Institutional Review Board, which determined that our study does not constitute human subjects research. The focus of the study is understanding website policies and practices, and emails associated with the study do not solicit personally identifiable information.

Note from Jonathan Mayer, the Principal Investigator (Saturday, December 18 @ 11:30pm)

Hi, my name is Jonathan Mayer. I'm the Principal Investigator for this academic research study. I have carefully read every single message sent to our research team, and I am dismayed that the emails in our study came across as security risks or legal threats. The intent of our study was to understand privacy practices, not to create a burden on website operators, email system operators, or privacy professionals. I sincerely apologize. I am the senior researcher, and the responsibility is mine.

The touchstone of my academic and government career, for over a decade, has been respecting and empowering users. That's why I study topics like web tracking, dark patterns, and broadband availability, and that's why I launched this study on privacy rights. I aim to be beyond reproach in my research methods, both out of principle and because my work often involves critiquing powerful companies and government agencies. In this instance, I fell short of that standard. I take your feedback to heart, and here is what I am doing about it.

Ethics

- Ethics is concerns the more general question about what kinds of conduct is considered “right” and/or “moral”

- Many traditions for deciding this, here are two of the most commonly invoked in the West:
  - **Consequentialism**
    - Focuses on outcomes – what are the consequences, positive and negative, of action
  - **Deontological ethics**
    - Focuses on whether action adherence to underlying norms/principles/moral duties; actions are fundamentally right or wrong, independent of outcomes

- What approach to use and how to apply it is not clear cut and different people have different opinions
  - Community norms change over time (sometimes significantly)
  - Zmap (next class) is a great example – from unethical to the norm

Aside: my experience

- Few people have a single consistent set of ethical principals
  - We apply different approaches in different circumstances or in combination
  - Our personal ethics are guided by experience and background
- For example
  - Much of our work has been guided by consequentialism, but there are lots of actions we would not take because we think the action itself is unethical and we undoubtedly have different standards of risk for criminal parties vs vulnerable parties
- The Kohno trolly problem paper is really about helping people surface what they think personally
- Different communities can feel quite differently about the same question
  - Note the differences between legal scholars, lawyers, ERB members and tech workers in the red lines paper – the potential *victims* were frequently the most comfortable with the research!

Re: Spamalytics

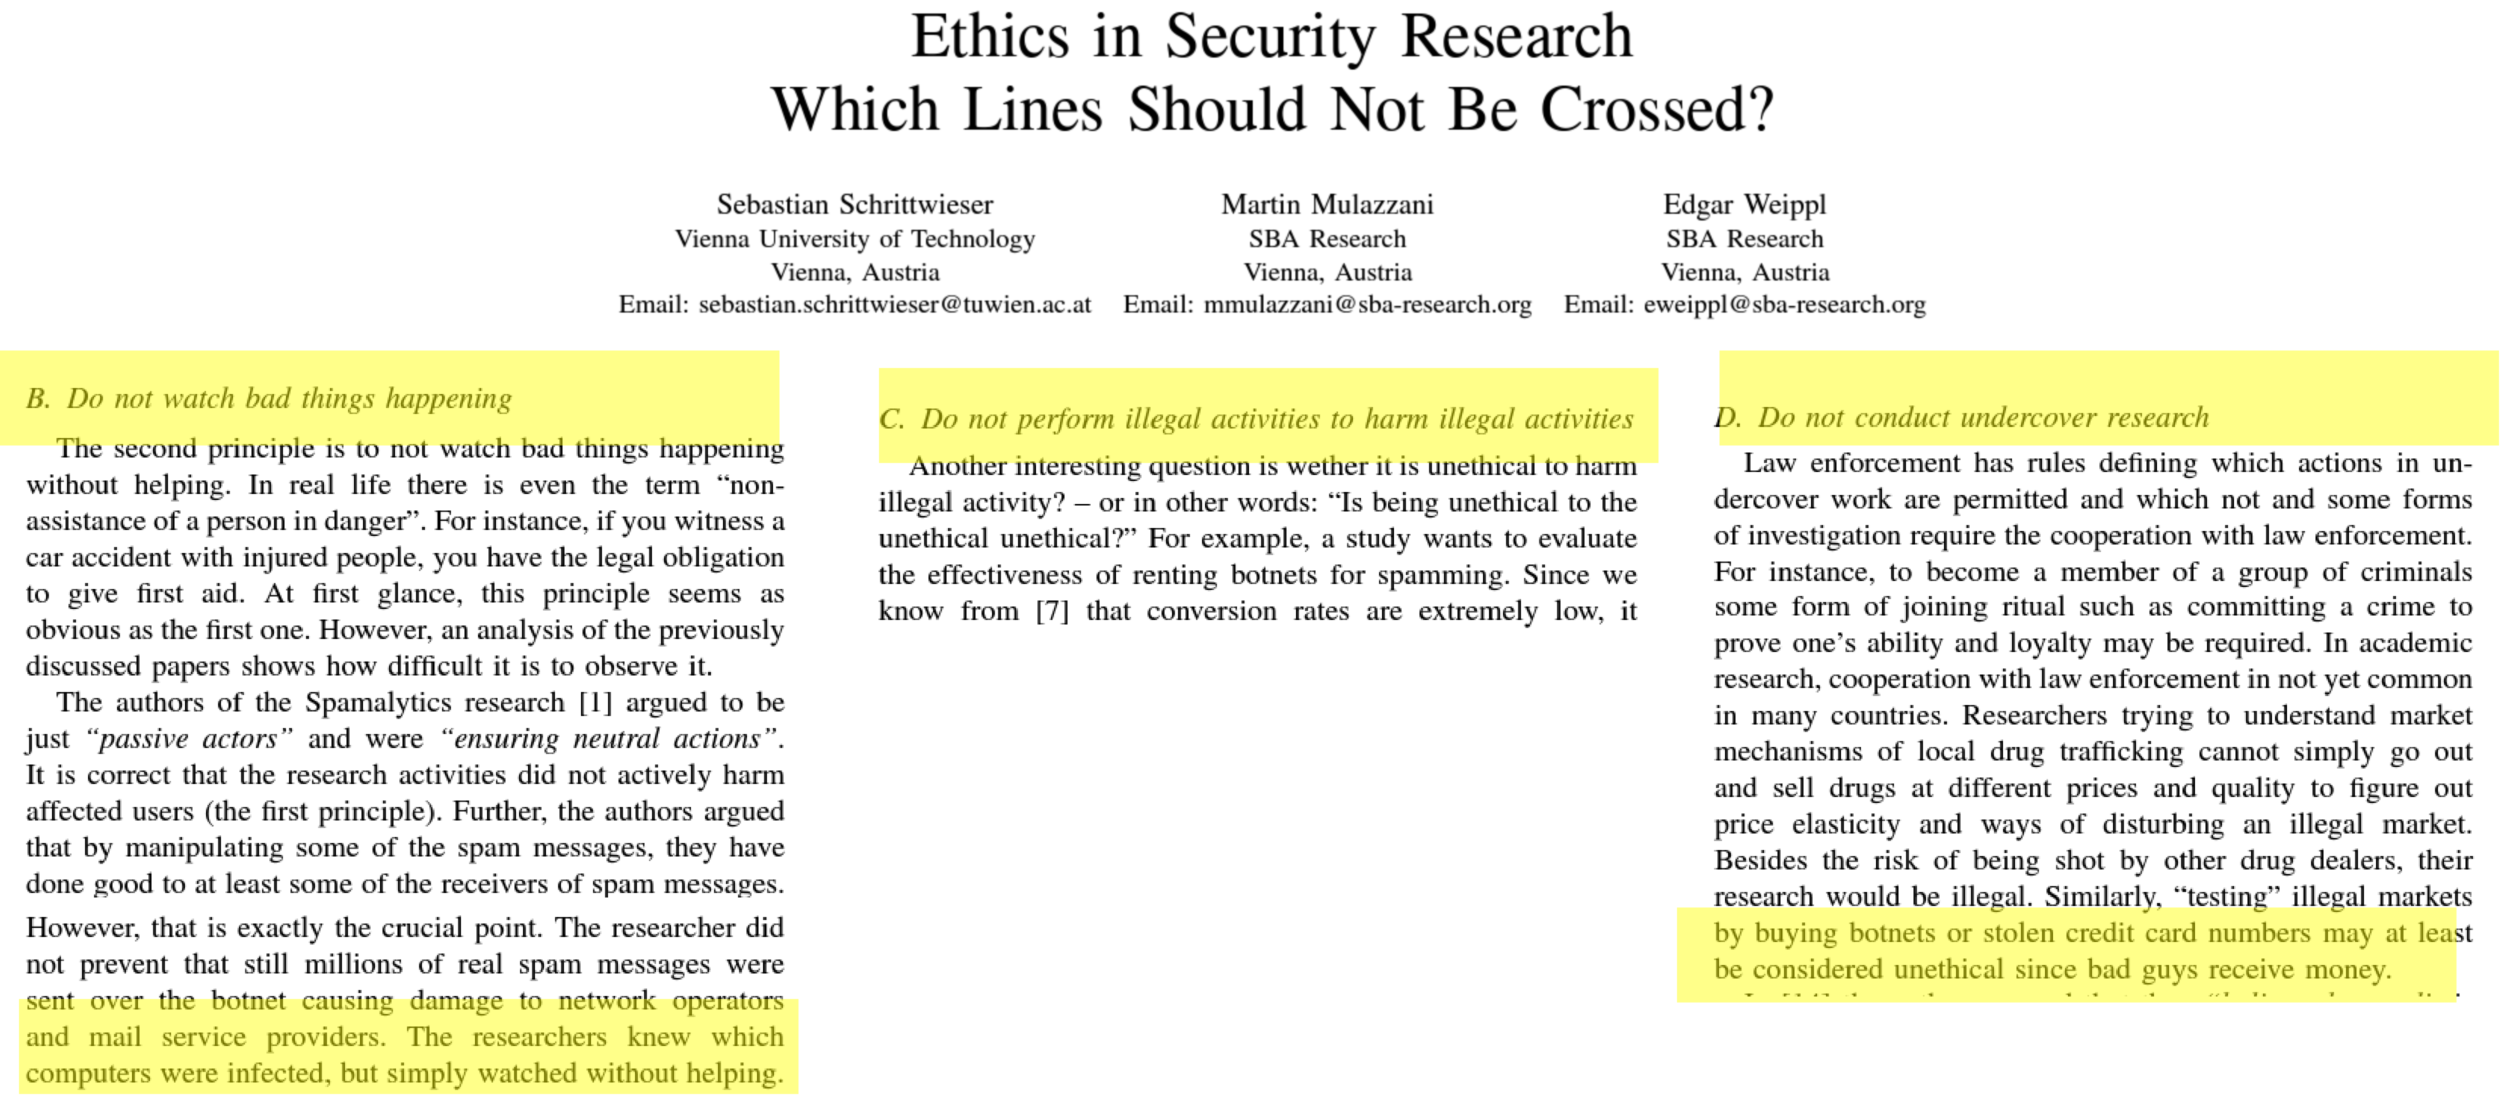

Ethics in Security Research
Which Lines Should Not Be Crossed?

Sebastian Schrittwieser
Vienna University of Technology
Vienna, Austria
Email: sebastian.schrittwieser@tuwien.ac.at

Martin Mulazzani
SBA Research
Vienna, Austria
Email: mmulazzani@sba-research.org

Edgar Weippl
SBA Research
Vienna, Austria
Email: eweippl@sba-research.org

B. Do not watch bad things happening

The second principle is to not watch bad things happening without helping. In real life there is even the term "non-assistance of a person in danger". For instance, if you witness a car accident with injured people, you have the legal obligation to give first aid. At first glance, this principle seems as obvious as the first one. However, an analysis of the previously discussed papers shows how difficult it is to observe it.

The authors of the Spamalytics research [1] argued to be just *"passive actors"* and were *"ensuring neutral actions"*. It is correct that the research activities did not actively harm affected users (the first principle). Further, the authors argued that by manipulating some of the spam messages, they have done good to at least some of the receivers of spam messages.

However, that is exactly the crucial point. The researcher did not prevent that still millions of real spam messages were sent over the botnet causing damage to network operators and mail service providers. The researchers knew which computers were infected, but simply watched without helping.

C. Do not perform illegal activities to harm illegal activities

Another interesting question is wether it is unethical to harm illegal activity? – or in other words: "Is being unethical to the unethical unethical?" For example, a study wants to evaluate the effectiveness of renting botnets for spamming. Since we know from [7] that conversion rates are extremely low, it

D. Do not conduct undercover research

Law enforcement has rules defining which actions in undercover work are permitted and which not and some forms of investigation require the cooperation with law enforcement. For instance, to become a member of a group of criminals some form of joining ritual such as committing a crime to prove one's ability and loyalty may be required. In academic research, cooperation with law enforcement in not yet common in many countries. Researchers trying to understand market mechanisms of local drug trafficking cannot simply go out and sell drugs at different prices and quality to figure out price elasticity and ways of disturbing an illegal market. Besides the risk of being shot by other drug dealers, their research would be illegal. Similarly, "testing" illegal markets by buying botnets or stolen credit card numbers may at least be considered unethical since bad guys receive money.

Re: Spamalytics

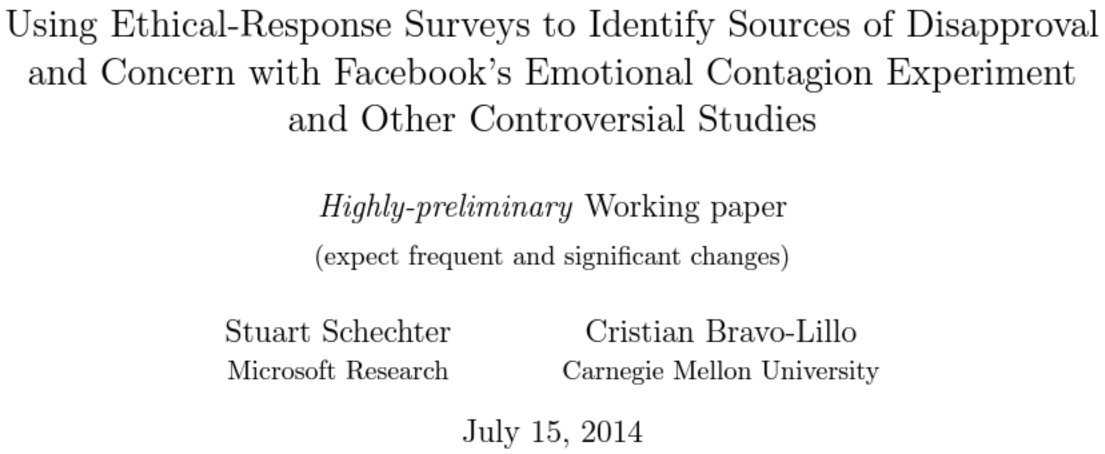

Using Ethical-Response Surveys to Identify Sources of Disapproval and Concern with Facebook's Emotional Contagion Experiment and Other Controversial Studies

Highly-preliminary Working paper

(expect frequent and significant changes)

Stuart Schechter
Microsoft Research

Cristian Bravo-Lillo
Carnegie Mellon University

July 15, 2014

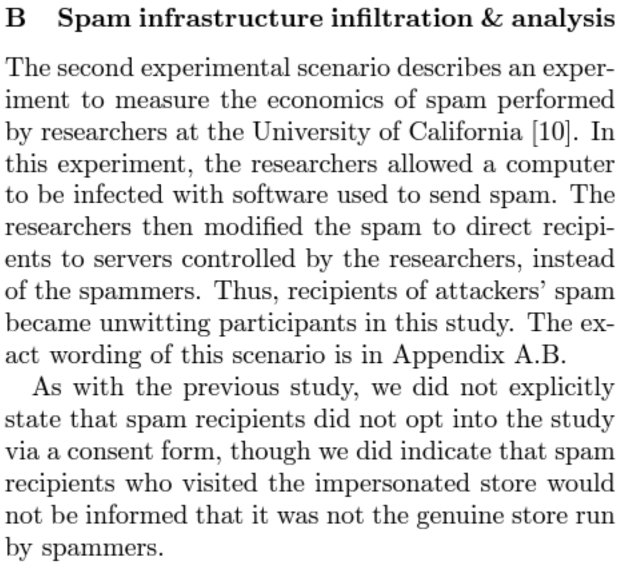

B Spam infrastructure infiltration & analysis

The second experimental scenario describes an experiment to measure the economics of spam performed by researchers at the University of California [10]. In this experiment, the researchers allowed a computer to be infected with software used to send spam. The researchers then modified the spam to direct recipients to servers controlled by the researchers, instead of the spammers. Thus, recipients of attackers' spam became unwitting participants in this study. The exact wording of this scenario is in Appendix A.B.

As with the previous study, we did not explicitly state that spam recipients did not opt into the study via a consent form, though we did indicate that spam recipients who visited the impersonated store would not be informed that it was not the genuine store run by spammers.

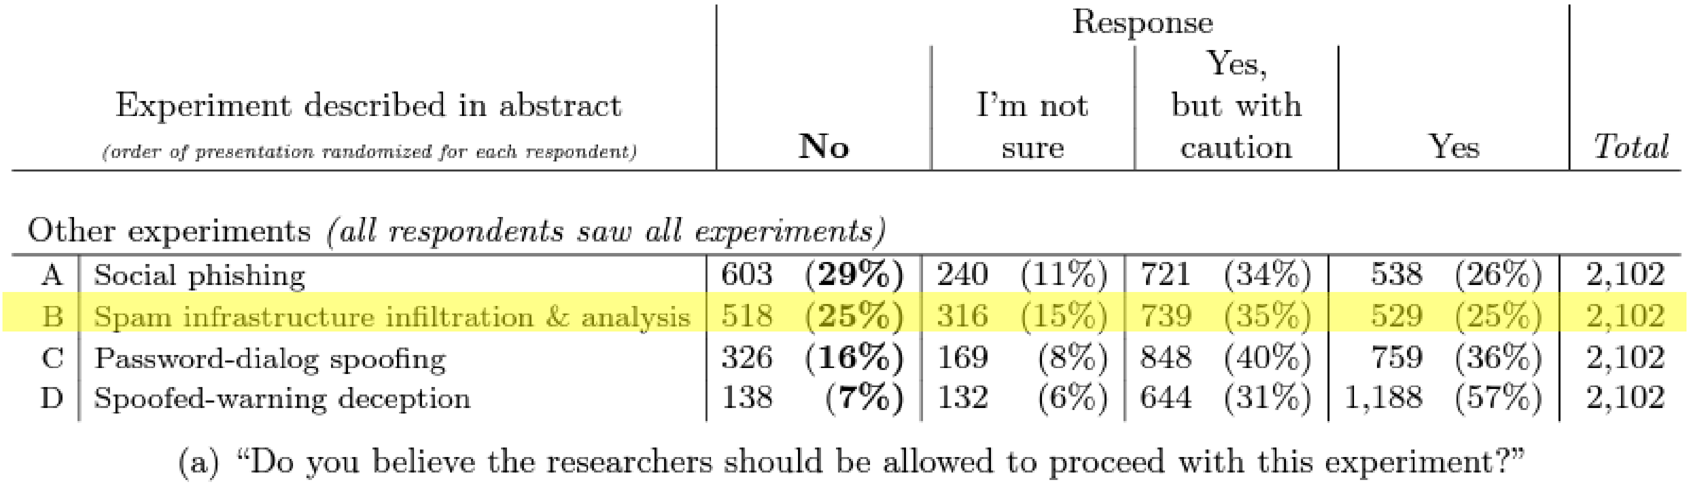

	Experiment described in abstract *(order of presentation randomized for each respondent)*	Response **No**		I'm not sure		Yes, but with caution		Yes		*Total*
	Other experiments *(all respondents saw all experiments)*									
A	Social phishing	603	**(29%)**	240	(11%)	721	(34%)	538	(26%)	2,102
B	Spam infrastructure infiltration & analysis	518	**(25%)**	316	(15%)	739	(35%)	529	(25%)	2,102
C	Password-dialog spoofing	326	**(16%)**	169	(8%)	848	(40%)	759	(36%)	2,102
D	Spoofed-warning deception	138	**(7%)**	132	(6%)	644	(31%)	1,188	(57%)	2,102

(a) "Do you believe the researchers should be allowed to proceed with this experiment?"

Taking a step back…

- What were the core ethical issues in each of those previous studies and what might have been inflammatory external factors?
- Facebook/Cornell Contagion study
  - Compare with previous Facebook/UCSD voting study
- Princeton CCPA study
  - Compare with UCSD study of F500 companies willingness to allow security review
- Hypocrite study
  - Compare with studies of embedded phishing training
  - Or studies that waste time of criminals
- Especially when research is about people, context can be huge factor
  - E.g. Geopolitical rivals < Criminals < Big Companies < Famous people < Normal people < Vulnerable people

Ethics not about humans

- How can there be ethical issues if the research isn’t about people?
  - Research not be about people, but may impact people
- Some examples:
  - Vulnerability disclosure: do you tell vendor before publication? Do you wait for them to fix issue? What if they can’t or don’t? What are your obligations to their customers?
  - What if your results will have significant legal/financial implications (i.e., you uncover a crypto blockchain is not solvent and so as a result there will be a run and investors will lose all their money)?
  - What if you end up outing a US govt counter-terrorism operation which may cause it to show down?
  - What if you do measurements that show people how to bypass TSA security
  - What if your actions (e.g., identifying how people evade censorship) will cause censorship to be improved/tightened?
- How far does one go in predicting the future? How far is reasonable?

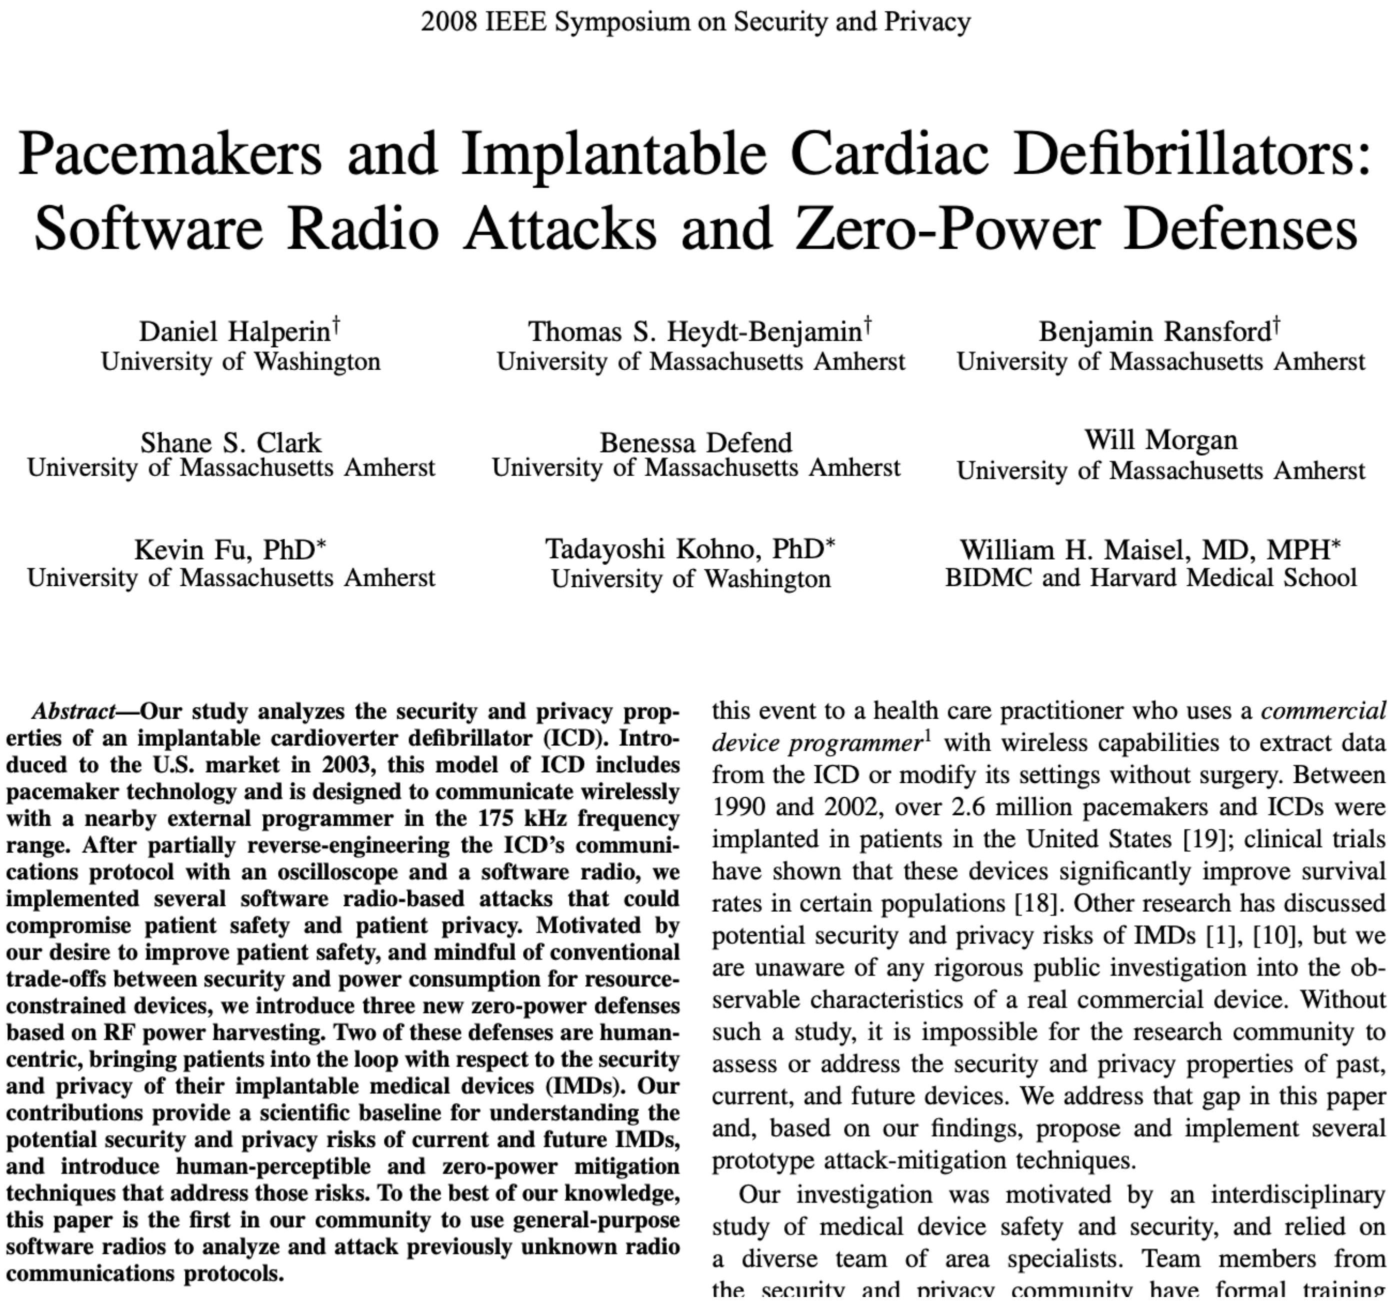

2008 IEEE Symposium on Security and Privacy

Pacemakers and Implantable Cardiac Defibrillators: Software Radio Attacks and Zero-Power Defenses

Daniel Halperin†
University of Washington

Thomas S. Heydt-Benjamin†
University of Massachusetts Amherst

Benjamin Ransford†
University of Massachusetts Amherst

Shane S. Clark
University of Massachusetts Amherst

Benessa Defend
University of Massachusetts Amherst

Will Morgan
University of Massachusetts Amherst

Kevin Fu, PhD*
University of Massachusetts Amherst

Tadayoshi Kohno, PhD*
University of Washington

William H. Maisel, MD, MPH*
BIDMC and Harvard Medical School

***Abstract*—Our study analyzes the security and privacy properties of an implantable cardioverter defibrillator (ICD). Introduced to the U.S. market in 2003, this model of ICD includes pacemaker technology and is designed to communicate wirelessly with a nearby external programmer in the 175 kHz frequency range. After partially reverse-engineering the ICD's communications protocol with an oscilloscope and a software radio, we implemented several software radio-based attacks that could compromise patient safety and patient privacy. Motivated by our desire to improve patient safety, and mindful of conventional trade-offs between security and power consumption for resource-constrained devices, we introduce three new zero-power defenses based on RF power harvesting. Two of these defenses are human-centric, bringing patients into the loop with respect to the security and privacy of their implantable medical devices (IMDs). Our contributions provide a scientific baseline for understanding the potential security and privacy risks of current and future IMDs, and introduce human-perceptible and zero-power mitigation techniques that address those risks. To the best of our knowledge, this paper is the first in our community to use general-purpose software radios to analyze and attack previously unknown radio communications protocols.**

this event to a health care practitioner who uses a *commercial device programmer*[1] with wireless capabilities to extract data from the ICD or modify its settings without surgery. Between 1990 and 2002, over 2.6 million pacemakers and ICDs were implanted in patients in the United States [19]; clinical trials have shown that these devices significantly improve survival rates in certain populations [18]. Other research has discussed potential security and privacy risks of IMDs [1], [10], but we are unaware of any rigorous public investigation into the observable characteristics of a real commercial device. Without such a study, it is impossible for the research community to assess or address the security and privacy properties of past, current, and future devices. We address that gap in this paper and, based on our findings, propose and implement several prototype attack-mitigation techniques.

Our investigation was motivated by an interdisciplinary study of medical device safety and security, and relied on a diverse team of area specialists. Team members from the security and privacy community have formal training

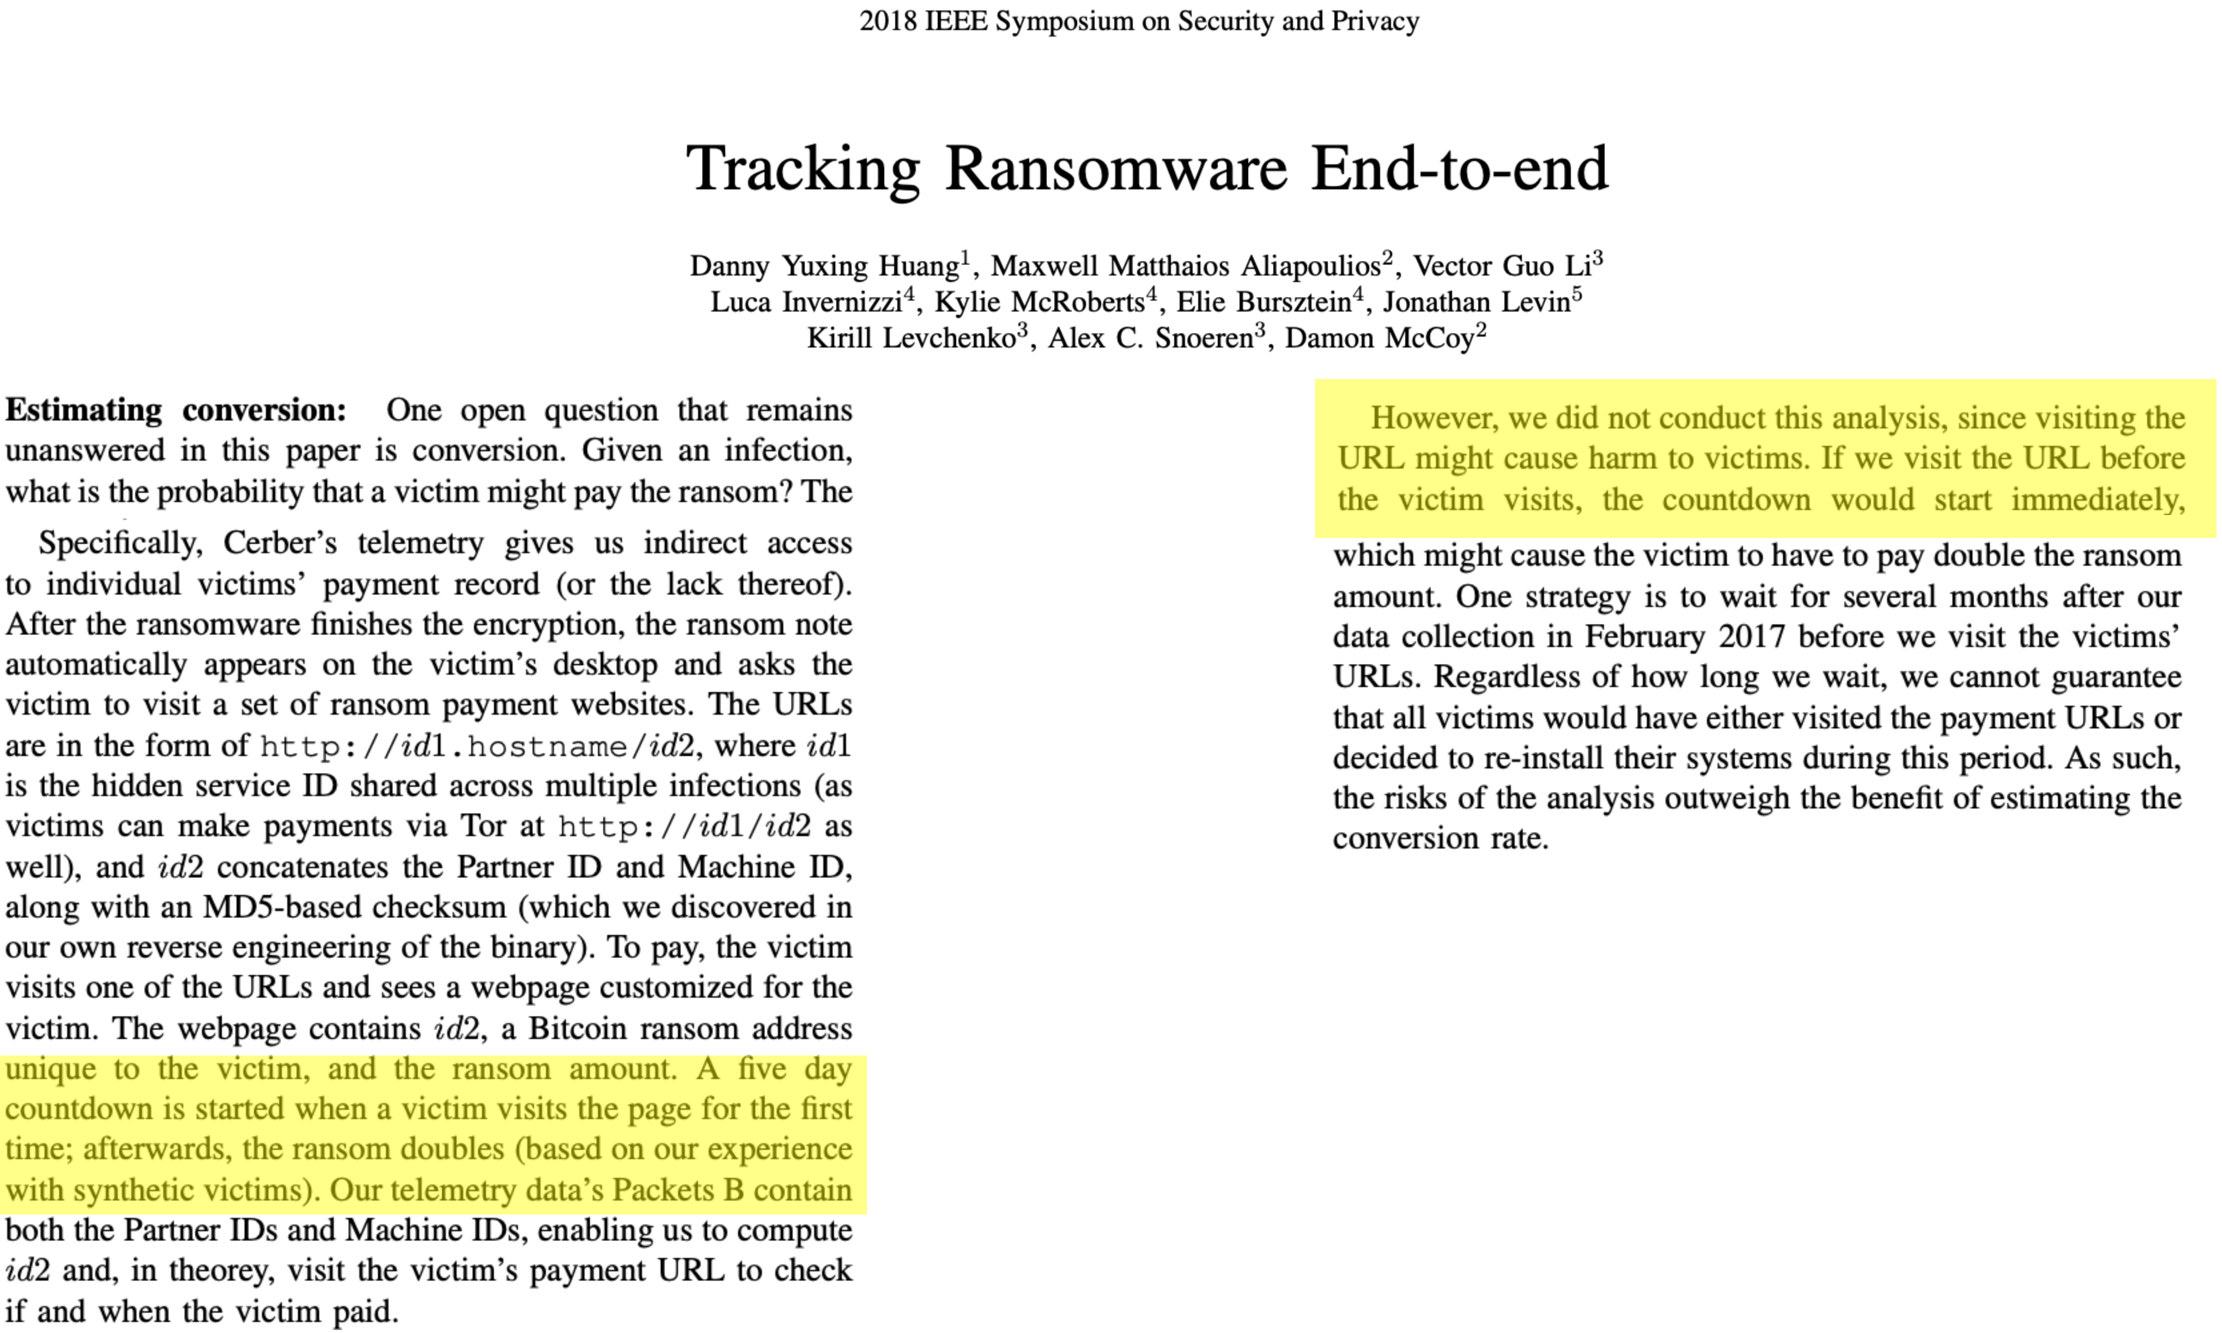

2018 IEEE Symposium on Security and Privacy

Tracking Ransomware End-to-end

Danny Yuxing Huang[1], Maxwell Matthaios Aliapoulios[2], Vector Guo Li[3]
Luca Invernizzi[4], Kylie McRoberts[4], Elie Bursztein[4], Jonathan Levin[5]
Kirill Levchenko[3], Alex C. Snoeren[3], Damon McCoy[2]

Estimating conversion: One open question that remains unanswered in this paper is conversion. Given an infection, what is the probability that a victim might pay the ransom? The

Specifically, Cerber's telemetry gives us indirect access to individual victims' payment record (or the lack thereof). After the ransomware finishes the encryption, the ransom note automatically appears on the victim's desktop and asks the victim to visit a set of ransom payment websites. The URLs are in the form of `http://id1.hostname/id2`, where $id1$ is the hidden service ID shared across multiple infections (as victims can make payments via Tor at `http://id1/id2` as well), and $id2$ concatenates the Partner ID and Machine ID, along with an MD5-based checksum (which we discovered in our own reverse engineering of the binary). To pay, the victim visits one of the URLs and sees a webpage customized for the victim. The webpage contains $id2$, a Bitcoin ransom address unique to the victim, and the ransom amount. A five day countdown is started when a victim visits the page for the first time; afterwards, the ransom doubles (based on our experience with synthetic victims). Our telemetry data's Packets B contain both the Partner IDs and Machine IDs, enabling us to compute $id2$ and, in theorey, visit the victim's payment URL to check if and when the victim paid.

However, we did not conduct this analysis, since visiting the URL might cause harm to victims. If we visit the URL before the victim visits, the countdown would start immediately, which might cause the victim to have to pay double the ransom amount. One strategy is to wait for several months after our data collection in February 2017 before we visit the victims' URLs. Regardless of how long we wait, we cannot guarantee that all victims would have either visited the payment URLs or decided to re-install their systems during this period. As such, the risks of the analysis outweigh the benefit of estimating the conversion rate.

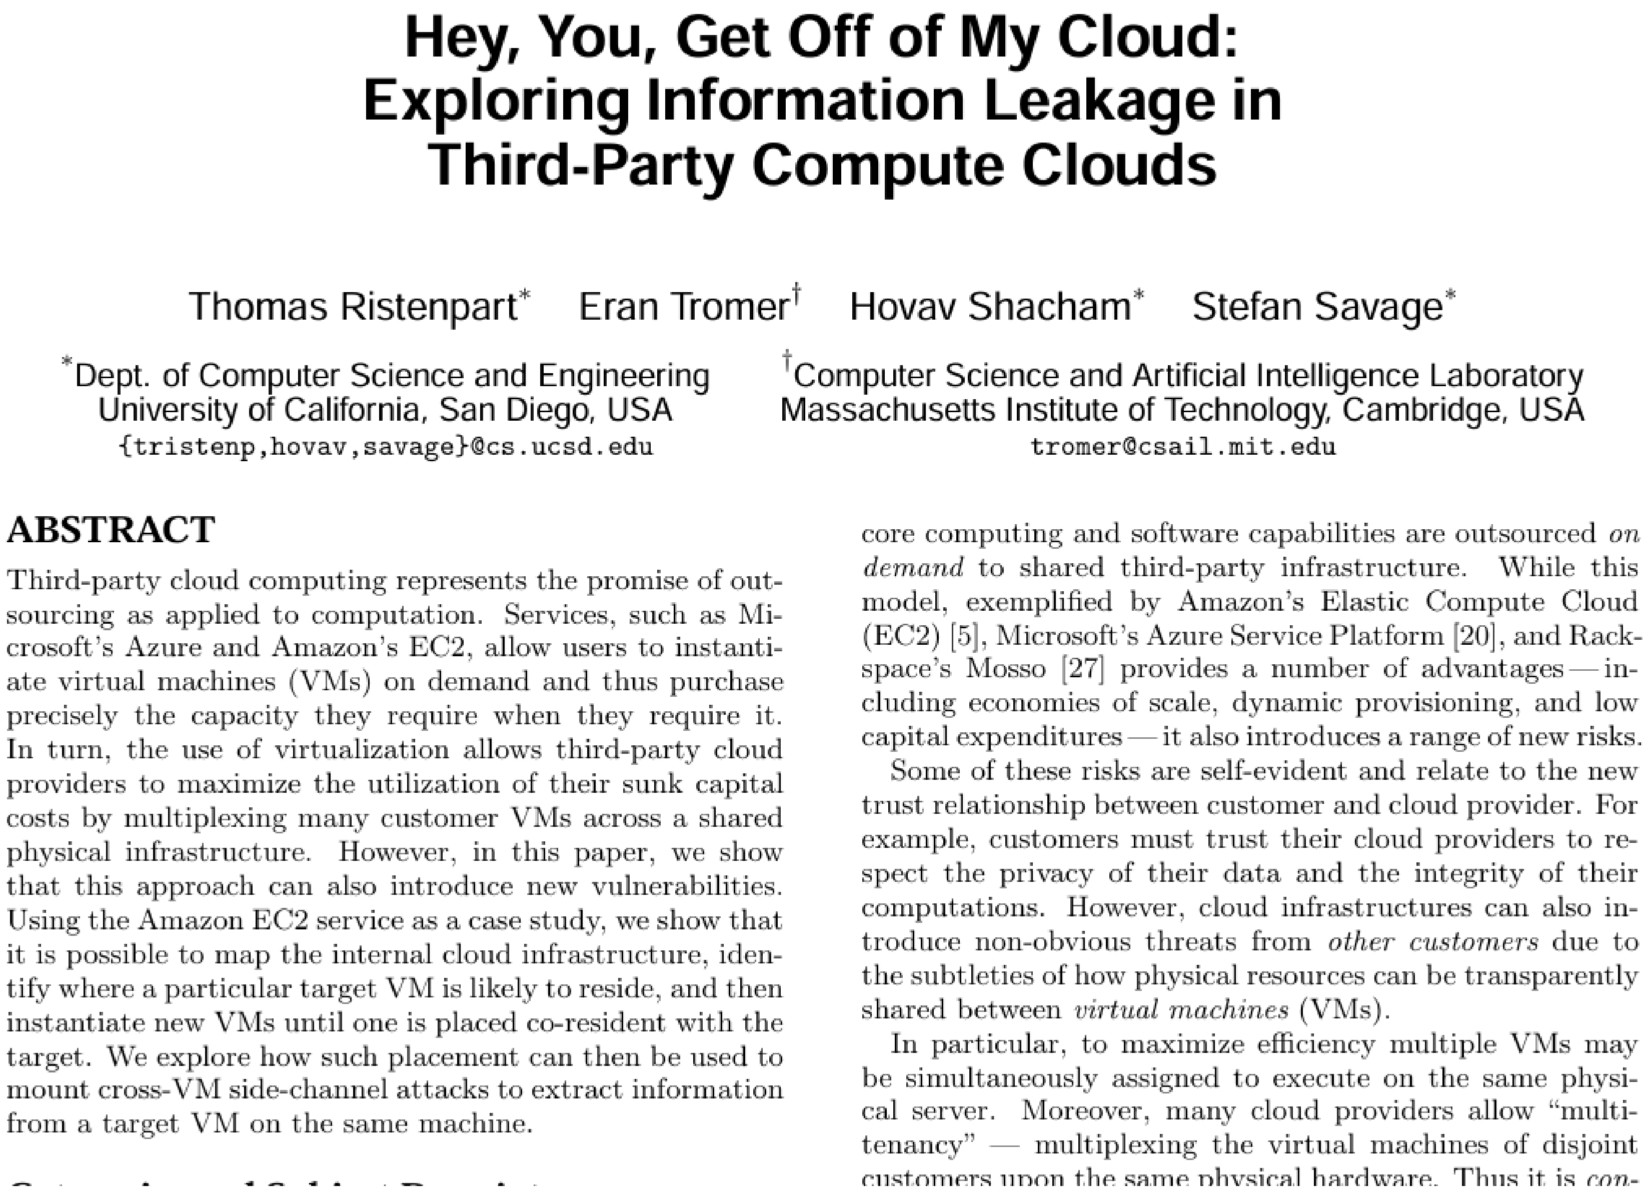

Hey, You, Get Off of My Cloud: Exploring Information Leakage in Third-Party Compute Clouds

Thomas Ristenpart* Eran Tromer† Hovav Shacham* Stefan Savage*

*Dept. of Computer Science and Engineering
University of California, San Diego, USA
{tristenp,hovav,savage}@cs.ucsd.edu

†Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology, Cambridge, USA
tromer@csail.mit.edu

ABSTRACT

Third-party cloud computing represents the promise of outsourcing as applied to computation. Services, such as Microsoft's Azure and Amazon's EC2, allow users to instantiate virtual machines (VMs) on demand and thus purchase precisely the capacity they require when they require it. In turn, the use of virtualization allows third-party cloud providers to maximize the utilization of their sunk capital costs by multiplexing many customer VMs across a shared physical infrastructure. However, in this paper, we show that this approach can also introduce new vulnerabilities. Using the Amazon EC2 service as a case study, we show that it is possible to map the internal cloud infrastructure, identify where a particular target VM is likely to reside, and then instantiate new VMs until one is placed co-resident with the target. We explore how such placement can then be used to mount cross-VM side-channel attacks to extract information from a target VM on the same machine.

Categories and Subject Descriptors

core computing and software capabilities are outsourced *on demand* to shared third-party infrastructure. While this model, exemplified by Amazon's Elastic Compute Cloud (EC2) [5], Microsoft's Azure Service Platform [20], and Rackspace's Mosso [27] provides a number of advantages — including economies of scale, dynamic provisioning, and low capital expenditures — it also introduces a range of new risks.

Some of these risks are self-evident and relate to the new trust relationship between customer and cloud provider. For example, customers must trust their cloud providers to respect the privacy of their data and the integrity of their computations. However, cloud infrastructures can also introduce non-obvious threats from *other customers* due to the subtleties of how physical resources can be transparently shared between *virtual machines* (VMs).

In particular, to maximize efficiency multiple VMs may be simultaneously assigned to execute on the same physical server. Moreover, many cloud providers allow "multi-tenancy" — multiplexing the virtual machines of disjoint customers upon the same physical hardware. Thus it is *con-*

Optics/complaints

- People can always be upset…
- Sometimes its for political reasons
  (e.g., disinformation research, climate science, origin of Covid, etc.)
- Sometimes its because they feel used/injured/angry
- Sometimes its because they have a unique take (or are misinformed) about legality, human subjects or ethics

- It is impossible to please all people all the time but if you can reasonably anticipate such a reaction (and this takes effort) you should think about your values and goals and what tradeoffs you might be willing to make

Questions you want to ask yourself when doing this work

- Am I doing anything that I can reasonably anticipate would upset someone?
- Can I change how I'm doing it to minimize that?
  - Add controls, get consent, anonymize data, don't call out entity by name, etc.
- How big is this risk and is the tradeoff worth it?
- If something were to go awry:
  - Have I done my due diligence in advance (i.e., will the university back you)
    e.g., received guidance from legal counsel, gotten IRB review, etc
  - Who is more sympathetic? You or the aggrieved party? Why is that?

Questions?

For next lecture

- Zmap