Secure Systems Engineering and Management

A Data-driven Approach

Measuring Direct and Indirect Outcomes

Michael Hicks
UPenn CIS 7000-003
Spring 2026

Evidence-based security

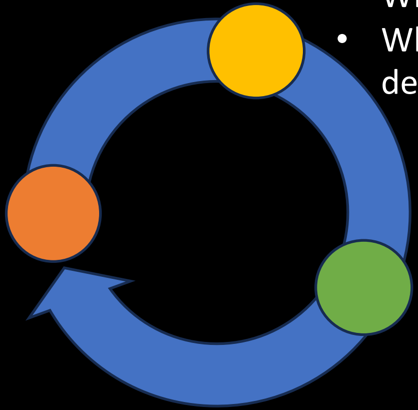

Analysis

- What are the (still) successful vectors of attack?
- Where is risk (still) greatest?
- What interventions could be deployed cost-effectively?

Intervention

- Engineering,
- Operations,
- Policy,
- Education, …

Outcomes

- Hosts compromised,
- Vulnerabilities exploited,
- Revenue lost, …

What to measure?

Outcomes

- Hosts compromised,
- Vulnerabilities exploited,
- Revenue lost, …

Intervention

- Engineering,
- Operations,
- Policy,
- Education, …

Measuring Security Practices and How They Impact Security

Louis F. DeKoven
University of California, San Diego
ldekoven@cs.ucsd.edu

Audrey Randall
University of California, San Diego
aurandal@eng.ucsd.edu

Ariana Mirian
University of California, San Diego
amirian@cs.ucsd.edu

Gautam Akiwate
University of California, San Diego
gakiwate@cs.ucsd.edu

Ansel Blume
University of California, San Diego
ablume@ucsd.edu

Lawrence K. Saul
University of California, San Diego
saul@cs.ucsd.edu

Aaron Schulman
University of California, San Diego
schulman@cs.ucsd.edu

Geoffrey M. Voelker
University of California, San Diego
voelker@cs.ucsd.edu

Stefan Savage
University of California, San Diego
savage@cs.ucsd.edu

ABSTRACT

Security is a discipline that places significant expectations on lay users. Thus, there are a wide array of technologies and behaviors that we exhort end users to adopt and thereby reduce their security risk. However, the adoption of these "best practices" — ranging from the use of antivirus products to actively keeping software updated — is not well understood, nor is their practical impact on security risk well-established. This paper explores both of these issues via a large-scale empirical measurement study covering approximately 15,000 computers over six months. We use passive monitoring to infer and characterize the prevalence of various security practices in situ as well as a range of other potentially security-relevant behaviors. We then explore the extent to which differences in key security behaviors impact real-world outcomes (i.e., that a device shows clear evidence of having been compromised).

CCS CONCEPTS

• **Security and privacy** → *Intrusion detection systems*; • **Networks**;

ACM Reference Format:
Louis F. DeKoven, Audrey Randall, Ariana Mirian, Gautam Akiwate, Ansel Blume, Lawrence K. Saul, Aaron Schulman, Geoffrey M. Voelker, and Stefan Savage. 2019. Measuring Security Practices and How They Impact Security. In *Internet Measurement Conference (IMC '19), October 21–23, 2019, Amsterdam, Netherlands.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3355369.3355571

1 INTRODUCTION

Ensuring effective computer security is widely understood to require a combination of both appropriate technological measures and prudent human behaviors; e.g., rapid installation of security updates to patch vulnerabilities or the use of password managers to ensure login credentials are distinct and random. Implicit in this status quo is the recognition that security is not an intrinsic property of today's systems, but is a byproduct of making appropriate choices — choices about what security products to employ, choices about how to manage system software, and choices about how to engage (or not) with third-party services on the Internet. Indeed, the codifying of good security choices, commonly referred to as security policy or "best practice", has been a part of our lives as long as security has been a concern.

However, establishing the value provided by these security practices is underexamined at best. First, we have limited empirical data about which security advice is adopted in practice. Users have a plethora of advice to choose from, highlighted by Reeder et al.'s recent study of expert security advice, whose title — "152 Simple Steps to Stay Safe Online" — underscores both the irony and the variability in such security lore [35]. Clearly few users are likely to follow all such dicta, but if user behavior is indeed key to security, it is important to know which practices are widely followed and which have only limited uptake.

A second, more subtle issue concerns the efficacy of security practices when followed: Do they work? Here the evidence is scant. Even practices widely agreed upon by Reeder's experts, such as keeping software patched, are not justified beyond a rhetorical argument. In fact, virtually all of the most established security best practices — including "use antivirus software", "use HTTPS/TLS", "update your software regularly", "use a password manager", and so on — have attained this status without empirical evidence quantifying their impact on security outcomes. Summarizing this state of affairs, Herley writes, "[Security] advice is complex and growing, but the benefit is largely speculative or moot", which he argues leads rational users to reject security advice [17].

To summarize, our existing models of security all rely on end users to follow a range of best practices. However, we neither understand the extent to which they are following this advice, nor do we have good information about how much this behavior ultimately impacts their future security.

This paper seeks to make progress on both issues — the prevalence of popular security practices and their relationship to security outcomes — via longitudinal empirical measurement of a large population of computer devices. In particular, we monitor the online

IMC '19, October 21–23, 2019, Amsterdam, Netherlands

ACM ISBN 978-1-4503-6948-0/19/10…$15.00
https://doi.org/10.1145/3355369.3355571

What to measure?

Outcomes

- Hosts compromised,
- Vulnerabilities exploited,
- Revenue lost, …

Intervention

- Engineering,
- Operations,
- Policy,
- Education, …

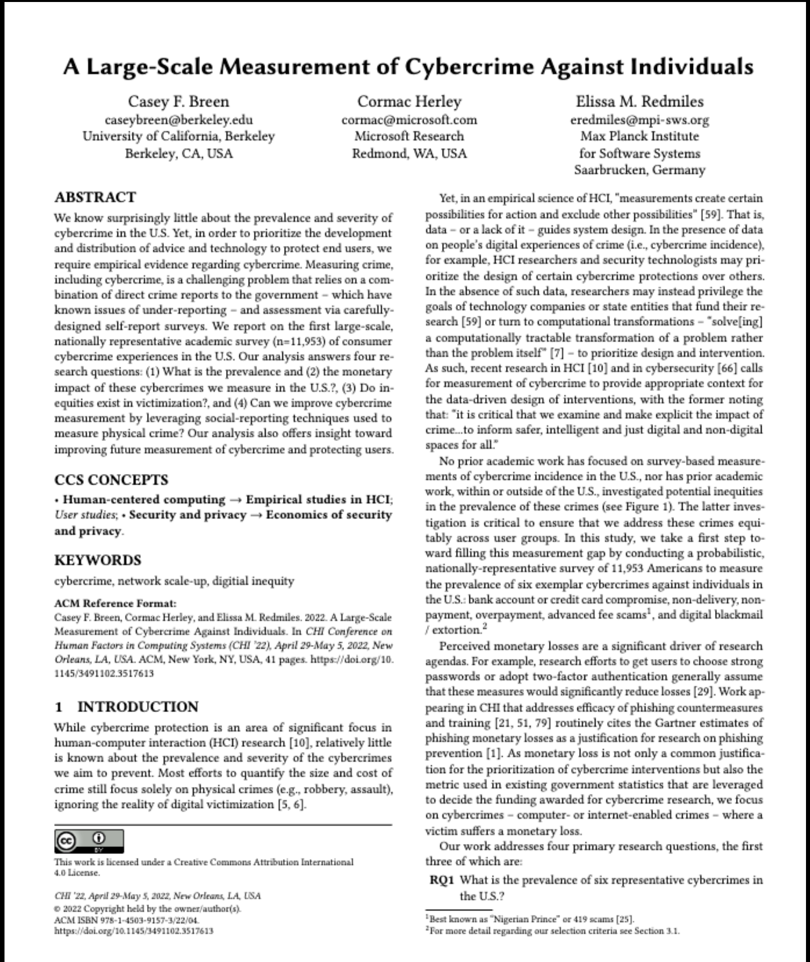

A Large-Scale Measurement of Cybercrime Against Individuals

Casey F. Breen
caseybreen@berkeley.edu
University of California, Berkeley
Berkeley, CA, USA

Cormac Herley
cormac@microsoft.com
Microsoft Research
Redmond, WA, USA

Elissa M. Redmiles
eredmiles@mpi-sws.org
Max Planck Institute
for Software Systems
Saarbrucken, Germany

ABSTRACT

We know surprisingly little about the prevalence and severity of cybercrime in the U.S. Yet, in order to prioritize the development and distribution of advice and technology to protect end users, we require empirical evidence regarding cybercrime. Measuring crime, including cybercrime, is a challenging problem that relies on a combination of direct crime reports to the government – which have known issues of under-reporting – and assessment via carefully-designed self-report surveys. We report on the first large-scale, nationally representative academic survey (n=11,953) of consumer cybercrime experiences in the U.S. Our analysis answers four research questions: (1) What is the prevalence and (2) the monetary impact of these cybercrimes we measure in the U.S.?, (3) Do inequities exist in victimization?, and (4) Can we improve cybercrime measurement by leveraging social-reporting techniques used to measure physical crime? Our analysis also offers insight toward improving future measurement of cybercrime and protecting users.

CCS CONCEPTS

• **Human-centered computing → Empirical studies in HCI**; *User studies*; • **Security and privacy → Economics of security and privacy**.

KEYWORDS

cybercrime, network scale-up, digitial inequity

ACM Reference Format:
Casey F. Breen, Cormac Herley, and Elissa M. Redmiles. 2022. A Large-Scale Measurement of Cybercrime Against Individuals. In *CHI Conference on Human Factors in Computing Systems (CHI '22), April 29-May 5, 2022, New Orleans, LA, USA*. ACM, New York, NY, USA, 41 pages. https://doi.org/10.1145/3491102.3517613

1 INTRODUCTION

While cybercrime protection is an area of significant focus in human-computer interaction (HCI) research [10], relatively little is known about the prevalence and severity of the cybercrimes we aim to prevent. Most efforts to quantify the size and cost of crime still focus solely on physical crimes (e.g., robbery, assault), ignoring the reality of digital victimization [5, 6].

Yet, in an empirical science of HCI, "measurements create certain possibilities for action and exclude other possibilities" [59]. That is, data – or a lack of it – guides system design. In the presence of data on people's digital experiences of crime (i.e., cybercrime incidence), for example, HCI researchers and security technologists may prioritize the design of certain cybercrime protections over others. In the absence of such data, researchers may instead privilege the goals of technology companies or state entities that fund their research [59] or turn to computational transformations – "solve[ing] a computationally tractable transformation of a problem rather than the problem itself" [7] – to prioritize design and intervention. As such, recent research in HCI [10] and in cybersecurity [66] calls for measurement of cybercrime to provide appropriate context for the data-driven design of interventions, with the former noting that: "it is critical that we examine and make explicit the impact of crime...to inform safer, intelligent and just digital and non-digital spaces for all."

No prior academic work has focused on survey-based measurements of cybercrime incidence in the U.S., nor has prior academic work, within or outside of the U.S., investigated potential inequities in the prevalence of these crimes (see Figure 1). The latter investigation is critical to ensure that we address these crimes equitably across user groups. In this study, we take a first step toward filling this measurement gap by conducting a probabilistic, nationally-representative survey of 11,953 Americans to measure the prevalence of six exemplar cybercrimes against individuals in the U.S.: bank account or credit card compromise, non-delivery, non-payment, overpayment, advanced fee scams[1], and digital blackmail / extortion.[2]

Perceived monetary losses are a significant driver of research agendas. For example, research efforts to get users to choose strong passwords or adopt two-factor authentication generally assume that these measures would significantly reduce losses [29]. Work appearing in CHI that addresses efficacy of phishing countermeasures and training [21, 51, 79] routinely cites the Gartner estimates of phishing monetary losses as a justification for research on phishing prevention [1]. As monetary loss is not only a common justification for the prioritization of cybercrime interventions but also the metric used in existing government statistics that are leveraged to decide the funding awarded for cybercrime research, we focus on cybercrimes – computer- or internet-enabled crimes – where a victim suffers a monetary loss.

Our work addresses four primary research questions, the first three of which are:

RQ1 What is the prevalence of six representative cybercrimes in the U.S.?

[1] Best known as "Nigerian Prince" or 419 scams [25].
[2] For more detail regarding our selection criteria see Section 3.1.

CHI '22, April 29-May 5, 2022, New Orleans, LA, USA
© 2022 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9157-3/22/04.
https://doi.org/10.1145/3491102.3517613

Two Papers, One Question: How Bad Is It?

	DeKoven et al.
Analogy	“Does the medicine work?”
Measures	Whether security practices reduce compromise
Approach	Passive network monitoring

Complementary: Network monitoring sees what users can’t report; surveys capture the human experience that network data cannot.

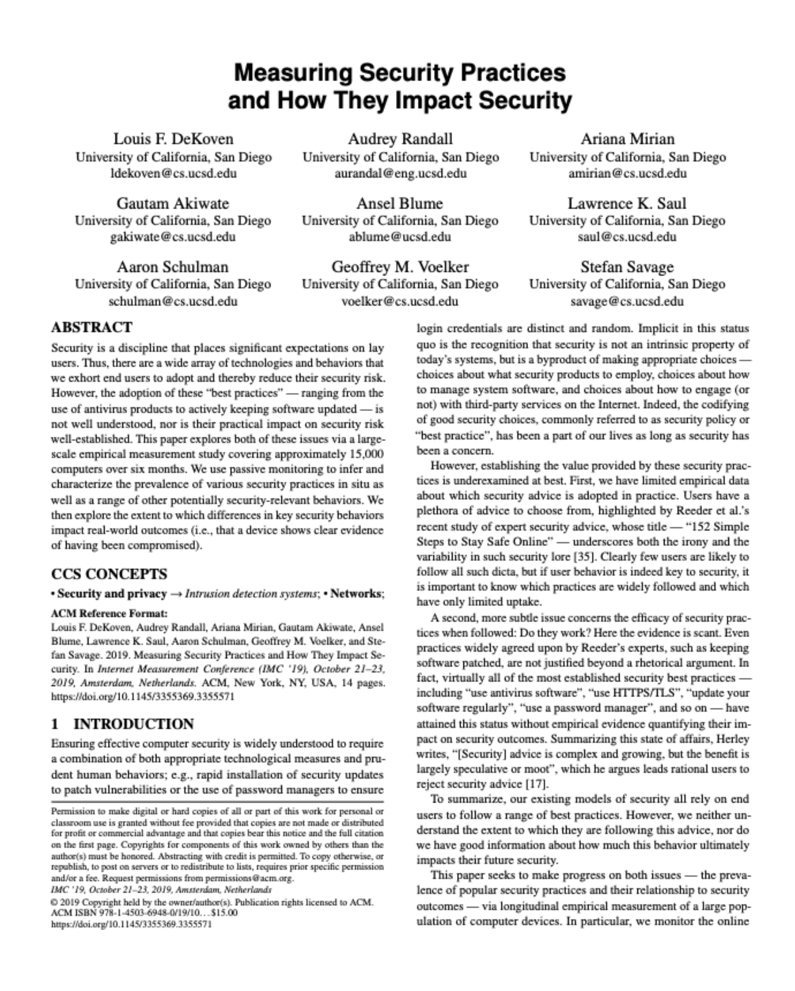

Measuring Security Practices and How They Impact Security

Louis F. DeKoven
University of California, San Diego
ldekoven@cs.ucsd.edu

Audrey Randall
University of California, San Diego
aurandal@eng.ucsd.edu

Ariana Mirian
University of California, San Diego
amirian@cs.ucsd.edu

Gautam Akiwate
University of California, San Diego
gakiwate@cs.ucsd.edu

Ansel Blume
University of California, San Diego
ablume@ucsd.edu

Lawrence K. Saul
University of California, San Diego
saul@cs.ucsd.edu

Aaron Schulman
University of California, San Diego
schulman@cs.ucsd.edu

Geoffrey M. Voelker
University of California, San Diego
voelker@cs.ucsd.edu

Stefan Savage
University of California, San Diego
savage@cs.ucsd.edu

ABSTRACT

Security is a discipline that places significant expectations on lay users. Thus, there are a wide array of technologies and behaviors that we exhort end users to adopt and thereby reduce their security risk. However, the adoption of these "best practices" — ranging from the use of antivirus products to actively keeping software updated — is not well understood, nor is their practical impact on security risk well-established. This paper explores both of these issues via a large-scale empirical measurement study covering approximately 15,000 computers over six months. We use passive monitoring to infer and characterize the prevalence of various security practices in situ as well as a range of other potentially security-relevant behaviors. We then explore the extent to which differences in key security behaviors impact real-world outcomes (i.e., that a device shows clear evidence of having been compromised).

CCS CONCEPTS

• **Security and privacy** → *Intrusion detection systems*; • **Networks**;

ACM Reference Format:
Louis F. DeKoven, Audrey Randall, Ariana Mirian, Gautam Akiwate, Ansel Blume, Lawrence K. Saul, Aaron Schulman, Geoffrey M. Voelker, and Stefan Savage. 2019. Measuring Security Practices and How They Impact Security. In *Internet Measurement Conference (IMC '19), October 21–23, 2019, Amsterdam, Netherlands.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3355369.3355571

1 INTRODUCTION

Ensuring effective computer security is widely understood to require a combination of both appropriate technological measures and prudent human behaviors; e.g., rapid installation of security updates to patch vulnerabilities or the use of password managers to ensure login credentials are distinct and random. Implicit in this status quo is the recognition that security is not an intrinsic property of today's systems, but is a byproduct of making appropriate choices — choices about what security products to employ, choices about how to manage system software, and choices about how to engage (or not) with third-party services on the Internet. Indeed, the codifying of good security choices, commonly referred to as security policy or "best practice", has been a part of our lives as long as security has been a concern.

However, establishing the value provided by these security practices is underexamined at best. First, we have limited empirical data about which security advice is adopted in practice. Users have a plethora of advice to choose from, highlighted by Reeder et al.'s recent study of expert security advice, whose title — "152 Simple Steps to Stay Safe Online" — underscores both the irony and the variability in such security lore [35]. Clearly few users are likely to follow all such dicta, but if user behavior is indeed key to security, it is important to know which practices are widely followed and which have only limited uptake.

A second, more subtle issue concerns the efficacy of security practices when followed: Do they work? Here the evidence is scant. Even practices widely agreed upon by Reeder's experts, such as keeping software patched, are not justified beyond a rhetorical argument. In fact, virtually all of the most established security best practices — including "use antivirus software", "use HTTPS/TLS", "update your software regularly", "use a password manager", and so on — have attained this status without empirical evidence quantifying their impact on security outcomes. Summarizing this state of affairs, Herley writes, "[Security] advice is complex and growing, but the benefit is largely speculative or moot", which he argues leads rational users to reject security advice [17].

To summarize, our existing models of security all rely on end users to follow a range of best practices. However, we neither understand the extent to which they are following this advice, nor do we have good information about how much this behavior ultimately impacts their future security.

This paper seeks to make progress on both issues — the prevalence of popular security practices and their relationship to security outcomes — via longitudinal empirical measurement of a large population of computer devices. In particular, we monitor the online

IMC '19, October 21–23, 2019, Amsterdam, Netherlands

ACM ISBN 978-1-4503-6948-0/19/10…$15.00
https://doi.org/10.1145/3355369.3355571

Research Questions and Security Outcomes

1. **Prevalence:** To what extent are common security "best practices" actually adopted by real users?

2. **Efficacy:** When adopted, do these practices actually reduce the likelihood of compromise?

Measurement Setup

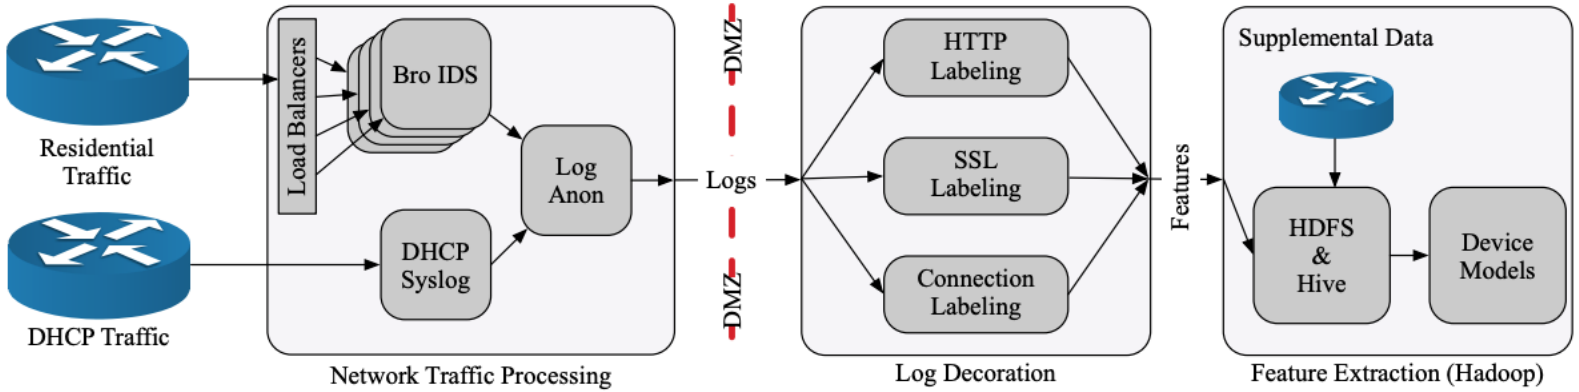

- **~15,000 desktop/laptop devices** on a university residential network
- **6 months** of observation (June–December 2018)

Inferring Device Features from Traffic

68 custom network signatures infer a device profile from traffic alone:

Traffic Signal	Inferred Feature
`Host: avast-update.com`	Antivirus: Avast
`User-Agent: … Windows NT 10.0`	OS: Windows 10
`DNS: tracker.torproject.org`	Tor usage
Version change in UA string	Software update event
IAB Content Taxonomy match	Website categories visited

Detecting Compromise

- Campus **Suricata IDS**
    - Restricted to **post-infection behavior** (e.g., C&C communication), not pre-compromise activity
- Rules manually curated to **remove frequent false-positive triggers**
- **Overall compromise rate: 4.5%** (682 of 15,291 devices)

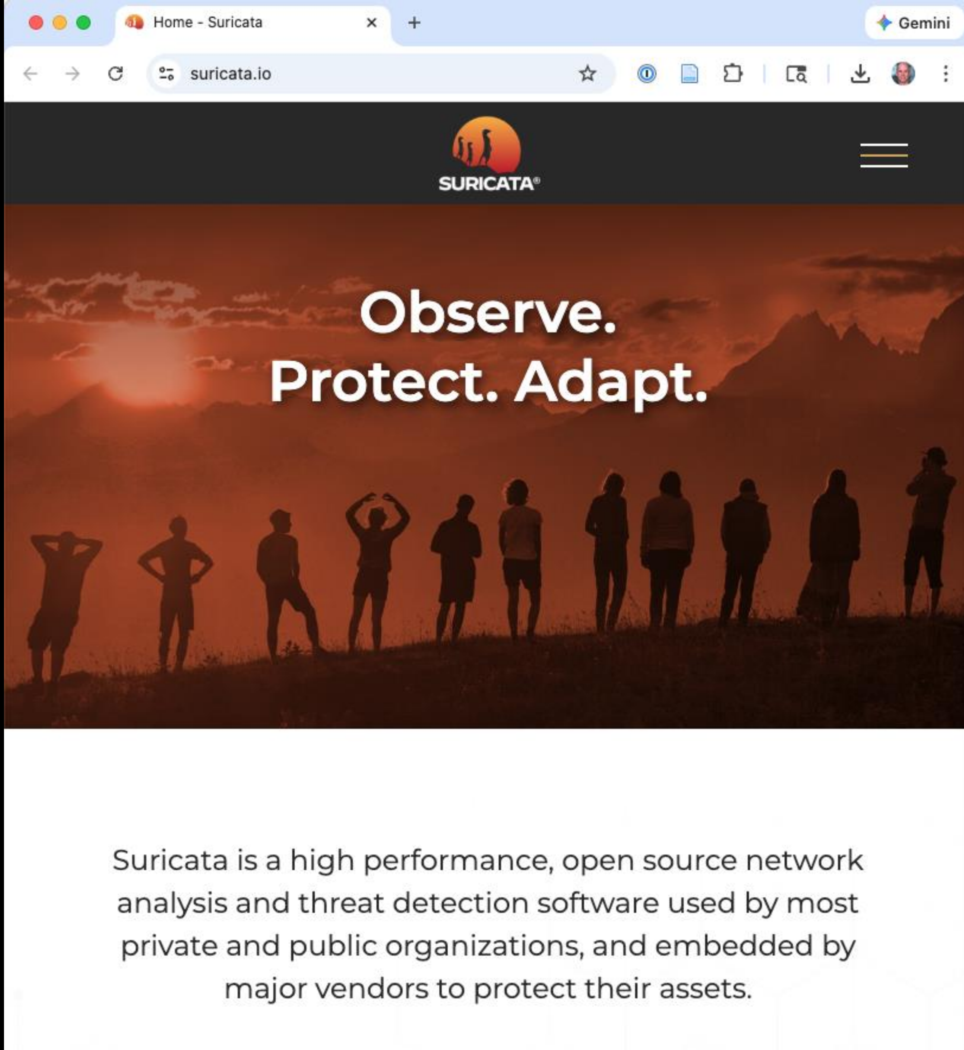

Observe. Protect. Adapt.

Suricata is a high performance, open source network analysis and threat detection software used by most private and public organizations, and embedded by major vendors to protect their assets.

Key Results: Individual Security Practices

Finding	Detail	Verdict
OS matters	Windows 3.9× more likely compromised than macOS	✅ Robust
Risky software	Thunderbird 33%, P2P 13%, Tor 12%, Adobe AIR 10%	✅ Robust
Web activity volume	More traffic → more compromise (“miles driven”)	✅ Robust
Gaming / hobby sites	Strongest content predictor of compromise	✅ Robust
Updates don’t help (much)	No significant protective effect; compromised devices updated *faster*	⚠ Surprising
AV doesn’t help (in this data)	AV users: 7% compromised vs. non-AV: 4%	⚠ Survivorship bias
Password managers	8% compromise (vs. 4% baseline)	⚠ Confounded by user profile

Statistical Methods Used

Test	Paper’s Use	What It Tests
Mann-Whitney U	OS/browser update speed	Two independent groups, same distribution?
Wilcoxon signed-rank	Before/after compromise updates	Paired observations differ?
Kolmogorov-Smirnov	Web content categories	Two distributions differ in any way?
Chi-Square	Software features vs. compromise	Association between categorical variables
Bonferroni	Applied to KS, Chi-Square	Controls family-wise error rate
LASSO logistic regression	Feature importance ranking (§6)	Binary classification + variable selection

Wilcoxon Signed-Rank Test

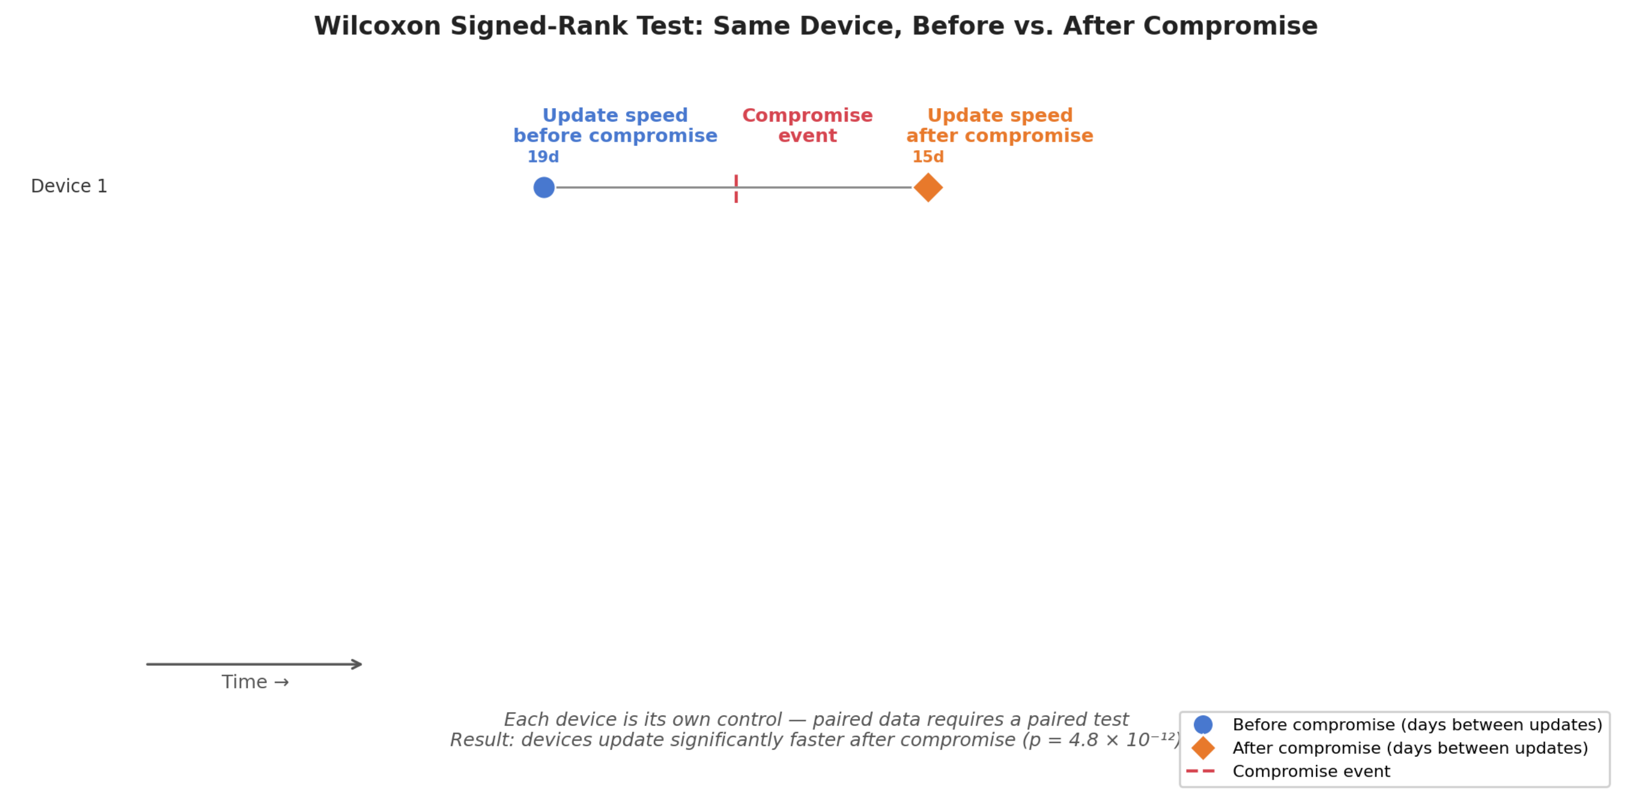

Wilcoxon Signed-Rank Test: Same Device, Before vs. After Compromise

Update speed before compromise
19d
Compromise event
Update speed after compromise
15d
Device 1
Time →
Before compromise (days between updates)
After compromise (days between updates)
Compromise event

Each device is its own control — paired data requires a paired test
Result: devices update significantly faster after compromise ($p = 4.8 \times 10^{-12}$)

Wilcoxon Signed-Rank Test

The paired-data cousin of Mann-Whitney U

- **Procedure:**
  1. Compute difference for each pair
  2. Rank the absolute differences
  3. Compare rank sums of positive vs. negative differences
- **Null hypothesis:** Differences are symmetric around zero

Paper's use: Devices update significantly faster *after* compromise than before ($p = 4.8 \times 10^{-12}$)

Which features matter most when considered together?

Approach: LASSO (L1 regularization) logistic regression

- **Why?**
  - Outcome is binary (compromised or not)
  - Hundreds of features → need automatic variable selection
    - Penalty shrinks unimportant coefficients to **exactly zero**
- **Why AUC-ROC?** With 4.5% base rate, accuracy is misleading
  - 95.5% by always predicting “clean”

Feature Ranking: Greedy Deletion

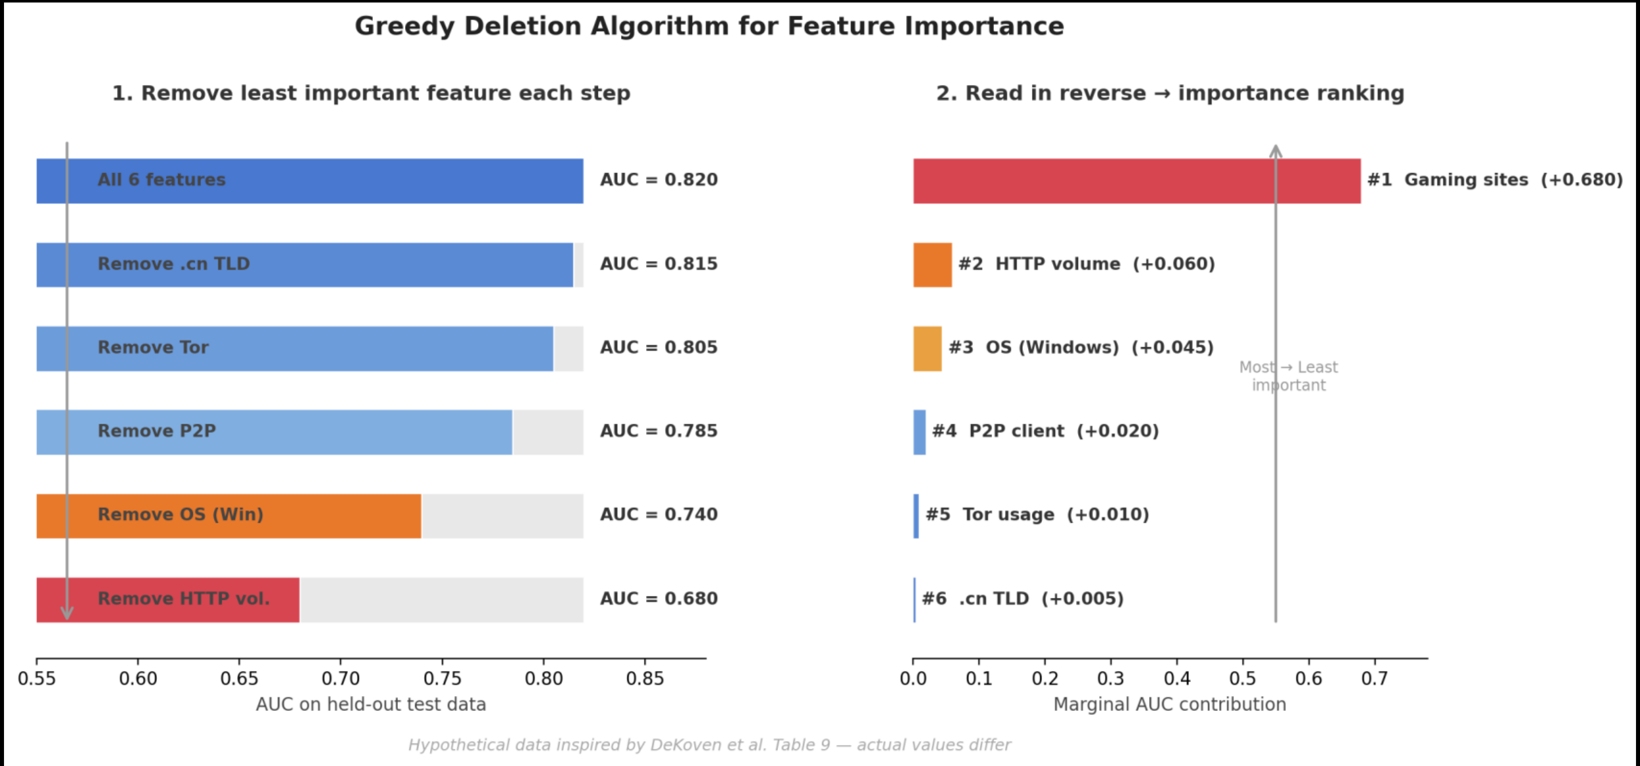

Hypothetical data inspired by DeKoven et al. Table 9 — actual values differ

Why Not Multiple Linear Regression with AIC?

	LASSO + Greedy Deletion
Complexity penalty	Aggressive; built-in variable selection
Evaluation	Out-of-sample performance (AUC)
Scalability	Handles hundreds of features well
Output	Clear feature ranking with AUC contributions

Key Results from LASSO Analysis

- **Behavioral features dominate** — web content and traffic volume matter more than security software
- **Visiting gaming sites** is the single strongest predictor of compromise (+68–76% AUC contribution)
- **HTTP traffic to .cn TLD** is a strong signal for macOS devices
- **IE User-Agent string** is an artifact — actually reflects QQ chat and Qihoo 360 usage, not Internet Explorer

Broader Lessons from DeKoven et al.

1. **Evidence-based security is hard** — confounders everywhere, randomized experiments nearly impossible
2. **Correlation ≠ causation is especially tricky** — Tor users are more compromised, but is Tor the cause?
3. **Best practices may signal more than they protect** — password manager users aren't compromised *because of* the tool
4. **Volume of activity is the strongest predictor** — more web activity → more exposure → more risk

A Large-Scale Measurement of Cybercrime Against Individuals

Casey F. Breen
caseybreen@berkeley.edu
University of California, Berkeley
Berkeley, CA, USA

Cormac Herley
cormac@microsoft.com
Microsoft Research
Redmond, WA, USA

Elissa M. Redmiles
eredmiles@mpi-sws.org
Max Planck Institute
for Software Systems
Saarbrucken, Germany

ABSTRACT

We know surprisingly little about the prevalence and severity of cybercrime in the U.S. Yet, in order to prioritize the development and distribution of advice and technology to protect end users, we require empirical evidence regarding cybercrime. Measuring crime, including cybercrime, is a challenging problem that relies on a combination of direct crime reports to the government – which have known issues of under-reporting – and assessment via carefully-designed self-report surveys. We report on the first large-scale, nationally representative academic survey (n=11,953) of consumer cybercrime experiences in the U.S. Our analysis answers four research questions: (1) What is the prevalence and (2) the monetary impact of these cybercrimes we measure in the U.S.?, (3) Do inequities exist in victimization?, and (4) Can we improve cybercrime measurement by leveraging social-reporting techniques used to measure physical crime? Our analysis also offers insight toward improving future measurement of cybercrime and protecting users.

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in HCI**; *User studies*; • **Security and privacy** → **Economics of security and privacy**.

KEYWORDS

cybercrime, network scale-up, digitial inequity

ACM Reference Format:
Casey F. Breen, Cormac Herley, and Elissa M. Redmiles. 2022. A Large-Scale Measurement of Cybercrime Against Individuals. In *CHI Conference on Human Factors in Computing Systems (CHI '22), April 29-May 5, 2022, New Orleans, LA, USA.* ACM, New York, NY, USA, 41 pages. https://doi.org/10.1145/3491102.3517613

1 INTRODUCTION

While cybercrime protection is an area of significant focus in human-computer interaction (HCI) research [10], relatively little is known about the prevalence and severity of the cybercrimes we aim to prevent. Most efforts to quantify the size and cost of crime still focus solely on physical crimes (e.g., robbery, assault), ignoring the reality of digital victimization [5, 6].

CHI '22, April 29-May 5, 2022, New Orleans, LA, USA
© 2022 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9157-3/22/04.
https://doi.org/10.1145/3491102.3517613

Yet, in an empirical science of HCI, "measurements create certain possibilities for action and exclude other possibilities" [59]. That is, data – or a lack of it – guides system design. In the presence of data on people's digital experiences of crime (i.e., cybercrime incidence), for example, HCI researchers and security technologists may prioritize the design of certain cybercrime protections over others. In the absence of such data, researchers may instead privilege the goals of technology companies or state entities that fund their research [59] or turn to computational transformations – "solve[ing] a computationally tractable transformation of a problem rather than the problem itself" [7] – to prioritize design and intervention. As such, recent research in HCI [10] and in cybersecurity [66] calls for measurement of cybercrime to provide appropriate context for the data-driven design of interventions, with the former noting that: "it is critical that we examine and make explicit the impact of crime...to inform safer, intelligent and just digital and non-digital spaces for all."

No prior academic work has focused on survey-based measurements of cybercrime incidence in the U.S., nor has prior academic work, within or outside of the U.S., investigated potential inequities in the prevalence of these crimes (see Figure 1). The latter investigation is critical to ensure that we address these crimes equitably across user groups. In this study, we take a first step toward filling this measurement gap by conducting a probabilistic, nationally-representative survey of 11,953 Americans to measure the prevalence of six exemplar cybercrimes against individuals in the U.S.: bank account or credit card compromise, non-delivery, non-payment, overpayment, advanced fee scams[1], and digital blackmail / extortion.[2]

Perceived monetary losses are a significant driver of research agendas. For example, research efforts to get users to choose strong passwords or adopt two-factor authentication generally assume that these measures would significantly reduce losses [29]. Work appearing in CHI that addresses efficacy of phishing countermeasures and training [21, 51, 79] routinely cites the Gartner estimates of phishing monetary losses as a justification for research on phishing prevention [1]. As monetary loss is not only a common justification for the prioritization of cybercrime interventions but also the metric used in existing government statistics that are leveraged to decide the funding awarded for cybercrime research, we focus on cybercrimes – computer- or internet-enabled crimes – where a victim suffers a monetary loss.

Our work addresses four primary research questions, the first three of which are:

RQ1 What is the prevalence of six representative cybercrimes in the U.S.?

[1] Best known as "Nigerian Prince" or 419 scams [25].

[2] For more detail regarding our selection criteria see Section 3.1.

Four Research Questions

1. **Prevalence** — What is the prevalence of various cybercrimes against U.S. *individuals*?
2. **Monetary Impact** — What is the financial cost?
3. **Inequities** — Do demographic disparities exist in who is victimized?
4. **NSUM** — Can the network scale-up method improve cybercrime measurement?

What are past approaches and their problems?

Source:

FBI Internet Crime Complaint Center (IC3)

- Victims of crime volunteer information
- Last 5 years: 2000 per day

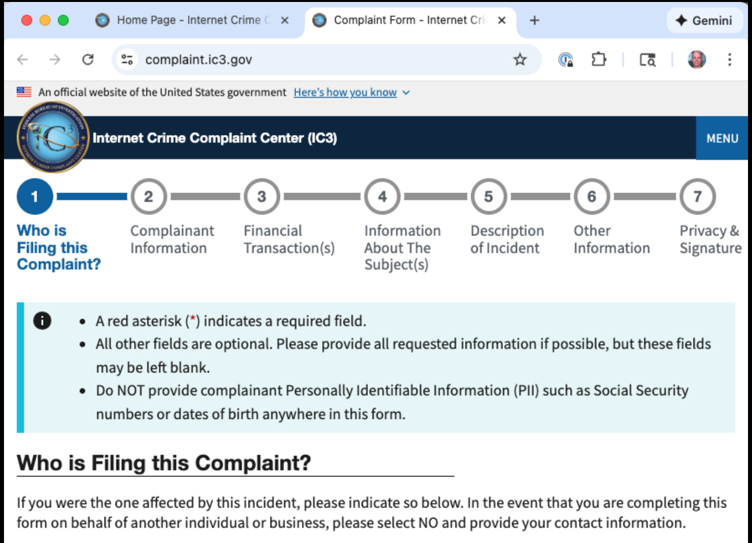

An official website of the United States government Here's how you know

Internet Crime Complaint Center (IC3) MENU

1 **Who is Filing this Complaint?** — 2 Complainant Information — 3 Financial Transaction(s) — 4 Information About The Subject(s) — 5 Description of Incident — 6 Other Information — 7 Privacy & Signature

- A red asterisk (*) indicates a required field.
- All other fields are optional. Please provide all requested information if possible, but these fields may be left blank.
- Do NOT provide complainant Personally Identifiable Information (PII) such as Social Security numbers or dates of birth anywhere in this form.

Who is Filing this Complaint?

If you were the one affected by this incident, please indicate so below. In the event that you are completing this form on behalf of another individual or business, please select NO and provide your contact information.

FBI IC3 estimates that only 10-12% of cybercrimes are reported

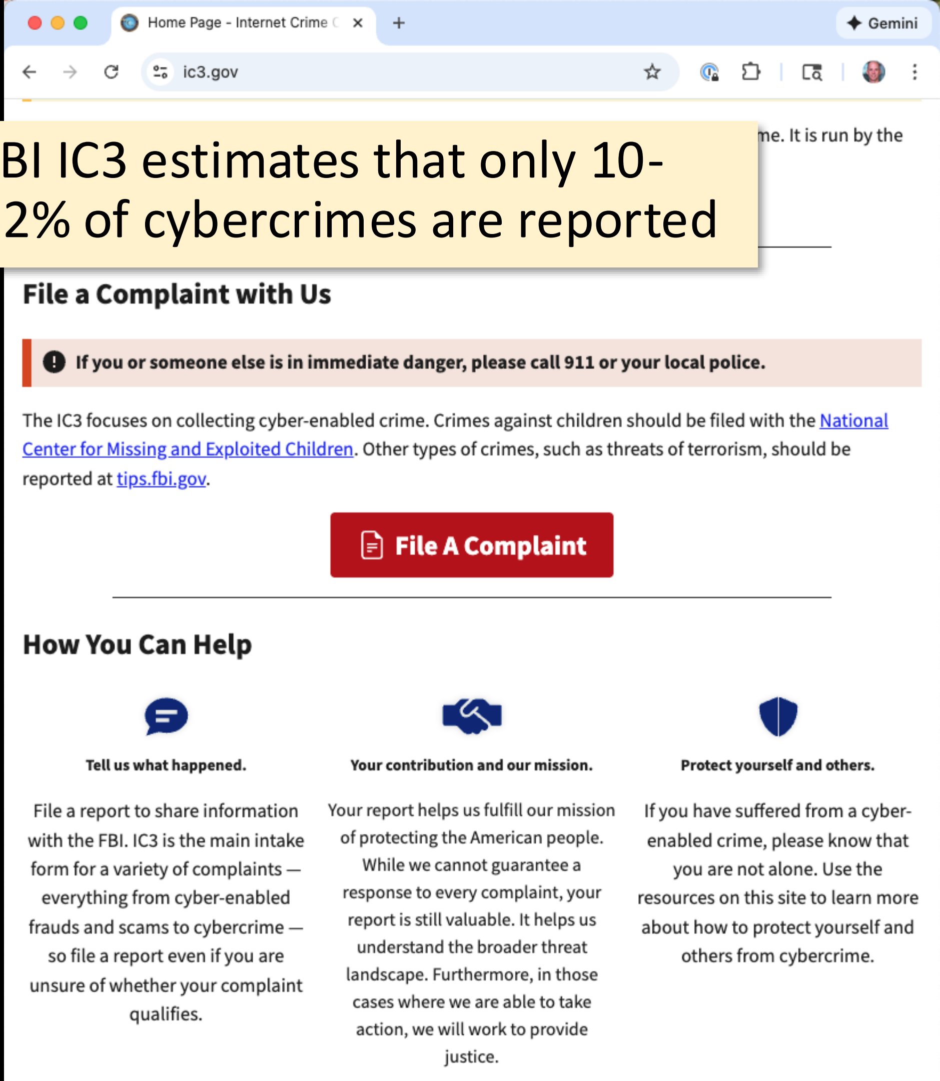

me. It is run by the

File a Complaint with Us

If you or someone else is in immediate danger, please call 911 or your local police.

The IC3 focuses on collecting cyber-enabled crime. Crimes against children should be filed with the National Center for Missing and Exploited Children. Other types of crimes, such as threats of terrorism, should be reported at tips.fbi.gov.

File A Complaint

How You Can Help

Tell us what happened.

File a report to share information with the FBI. IC3 is the main intake form for a variety of complaints — everything from cyber-enabled frauds and scams to cybercrime — so file a report even if you are unsure of whether your complaint qualifies.

Your contribution and our mission.

Your report helps us fulfill our mission of protecting the American people. While we cannot guarantee a response to every complaint, your report is still valuable. It helps us understand the broader threat landscape. Furthermore, in those cases where we are able to take action, we will work to provide justice.

Protect yourself and others.

If you have suffered from a cyber-enabled crime, please know that you are not alone. Use the resources on this site to learn more about how to protect yourself and others from cybercrime.

Source

Recruiting paid participants on-line

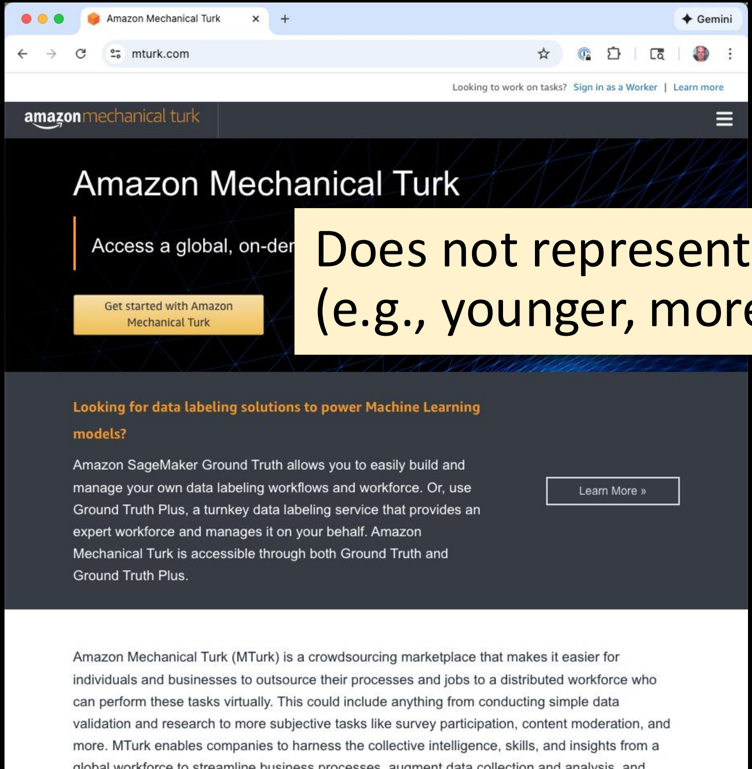

Looking to work on tasks? Sign in as a Worker | Learn more

amazon mechanical turk

Amazon Mechanical Turk

Access a global, on-der

Get started with Amazon Mechanical Turk

Looking for data labeling solutions to power Machine Learning models?

Amazon SageMaker Ground Truth allows you to easily build and manage your own data labeling workflows and workforce. Or, use Ground Truth Plus, a turnkey data labeling service that provides an expert workforce and manages it on your behalf. Amazon Mechanical Turk is accessible through both Ground Truth and Ground Truth Plus.

Learn More »

Amazon Mechanical Turk (MTurk) is a crowdsourcing marketplace that makes it easier for individuals and businesses to outsource their processes and jobs to a distributed workforce who can perform these tasks virtually. This could include anything from conducting simple data validation and research to more subjective tasks like survey participation, content moderation, and more. MTurk enables companies to harness the collective intelligence, skills, and insights from a global workforce to streamline business processes, augment data collection and analysis, and

Does not represent whole US population (e.g., younger, more Internet-savvy)

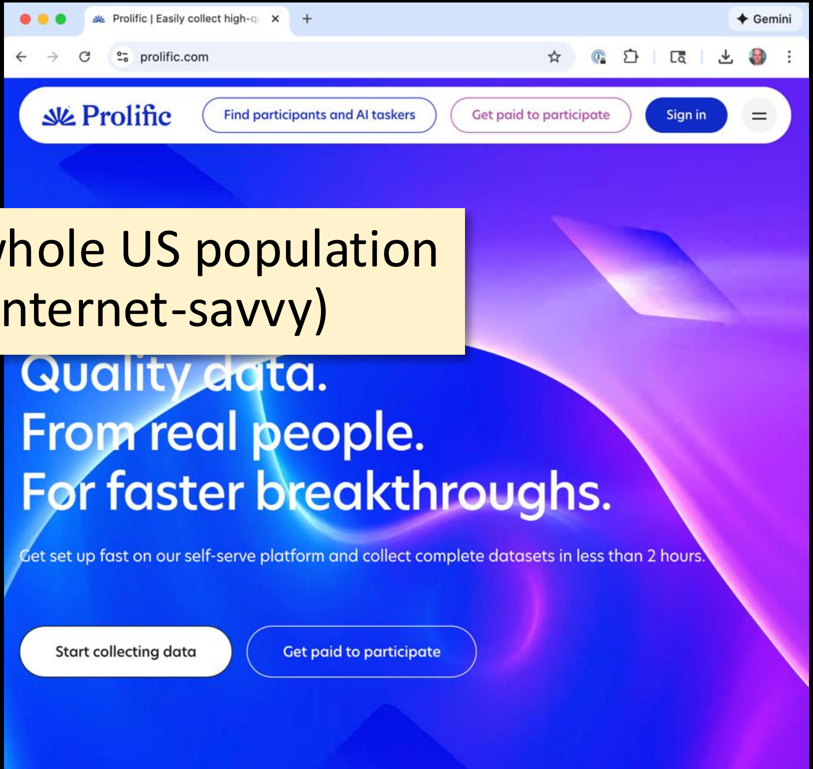
Prolific | Easily collect high-q
Gemini
prolific.com
Prolific
Find participants and AI taskers
Get paid to participate
Sign in

Quality data.
From real people.
For faster breakthroughs.

Get set up fast on our self-serve platform and collect complete datasets in less than 2 hours.

Start collecting data

Get paid to participate

Whole population

Multistage (and probabilistic) sampling

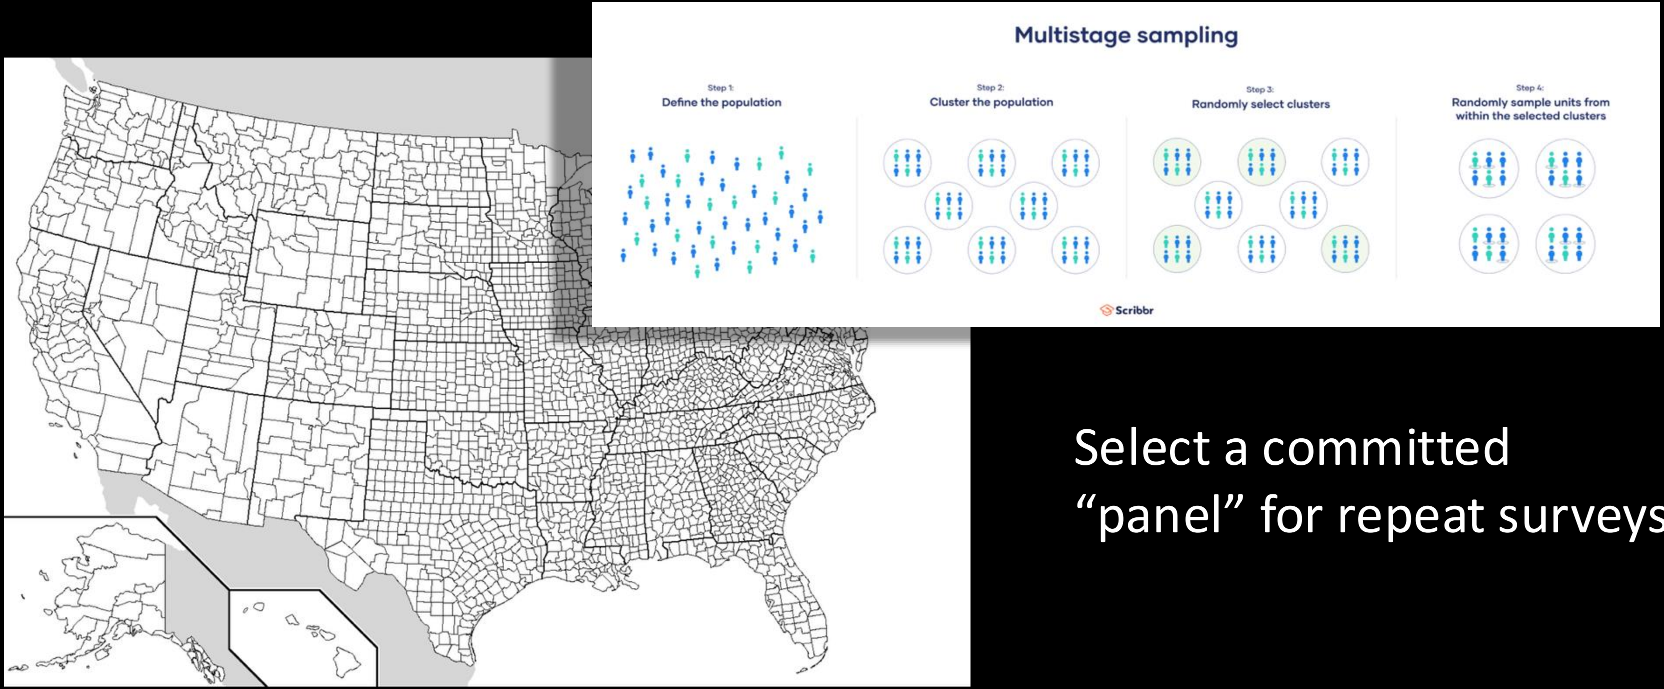

Select a committed “panel” for repeat surveys

AmeriSpeak Panel

- **Probability-based panel** operated by NORC at the University of Chicago
- Households selected via **area probability + address-based sampling**
- Multi-mode recruitment: mail, telephone, **in-person field interviewers**
- Covers ~**97% of U.S. household population**
- Used by major government-funded studies

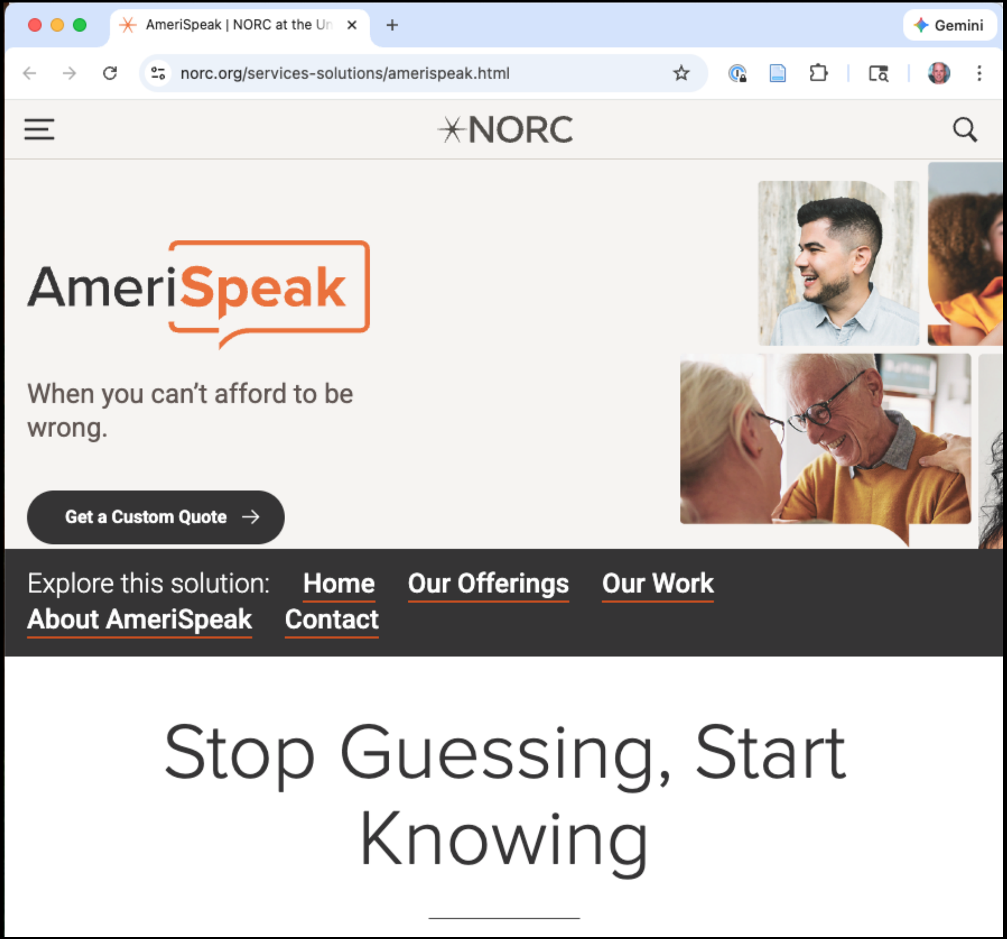

AmeriSpeak | NORC at the Un

norc.org/services-solutions/amerispeak.html

Gemini

NORC

AmeriSpeak

When you can't afford to be wrong.

Get a Custom Quote →

Explore this solution: Home Our Offerings Our Work About AmeriSpeak Contact

Stop Guessing, Start Knowing

Aside: Concern with response rate

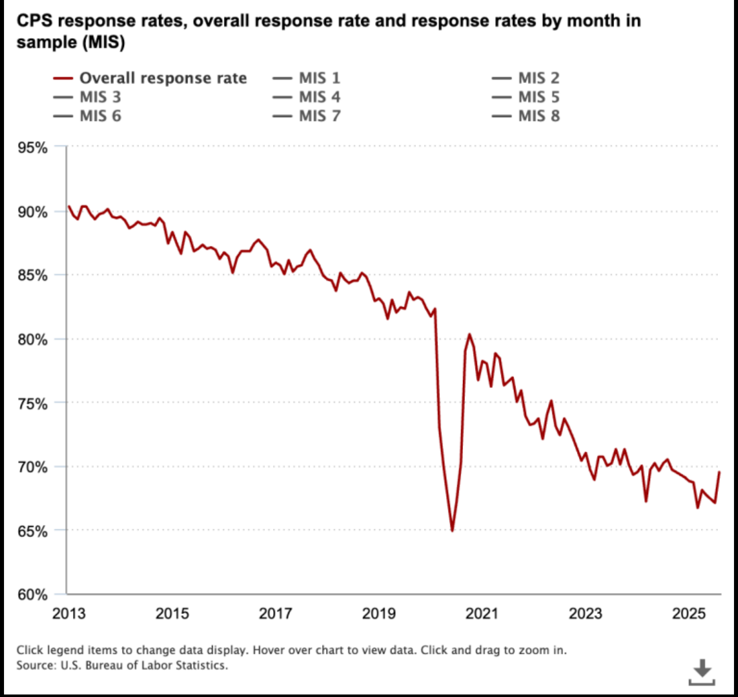

CPS response rates, overall response rate and response rates by month in sample (MIS)

Click legend items to change data display. Hover over chart to view data. Click and drag to zoom in.
Source: U.S. Bureau of Labor Statistics.

Two Samples, One Study

	General Sample	“Rare Event” Sample
N	1,002	10,951
Purpose	Estimate common crimes	Estimate rare crimes
Crime base rates	High enough (~6–23%)	Very low (<1%)

Why two sizes? Statistical power for rare events

- A 0.5% crime in N=1,002 → ~5 victims (too few)
- A 0.5% crime in N=10,951 → ~55 victims (enough for basic estimation)

RQ1: Cybercrime Prevalence

Roughly 1 in 8 U.S. adults per year.

Rates are **much higher** than FBI IC3 data

Crime Type	Annualized Prevalence
Banking / Credit Card Fraud	~12%
Non-Delivery Scam	~3%
Non-Payment	~0.3%
Overpayment	~0.2%
Extortion	~0.1%
Advance Fee	~0.1%

Based on the general sample

How Probability Sampling Changes Statistics

Without survey weights (e.g., DeKoven et al.):

- Proportion = count / total
- Standard CI formulas apply directly

With survey weights (Breen et al.):

- Weighted proportion = Σ(weight × indicator) / Σ(weight)
- CI must account for survey design (not just sample size)

Source:

National Crime Victimization Survey

- Annual interview of representative sample of about 240,000 persons in about 150,000 US households
- Non-fatal personal crimes, officially reported or not

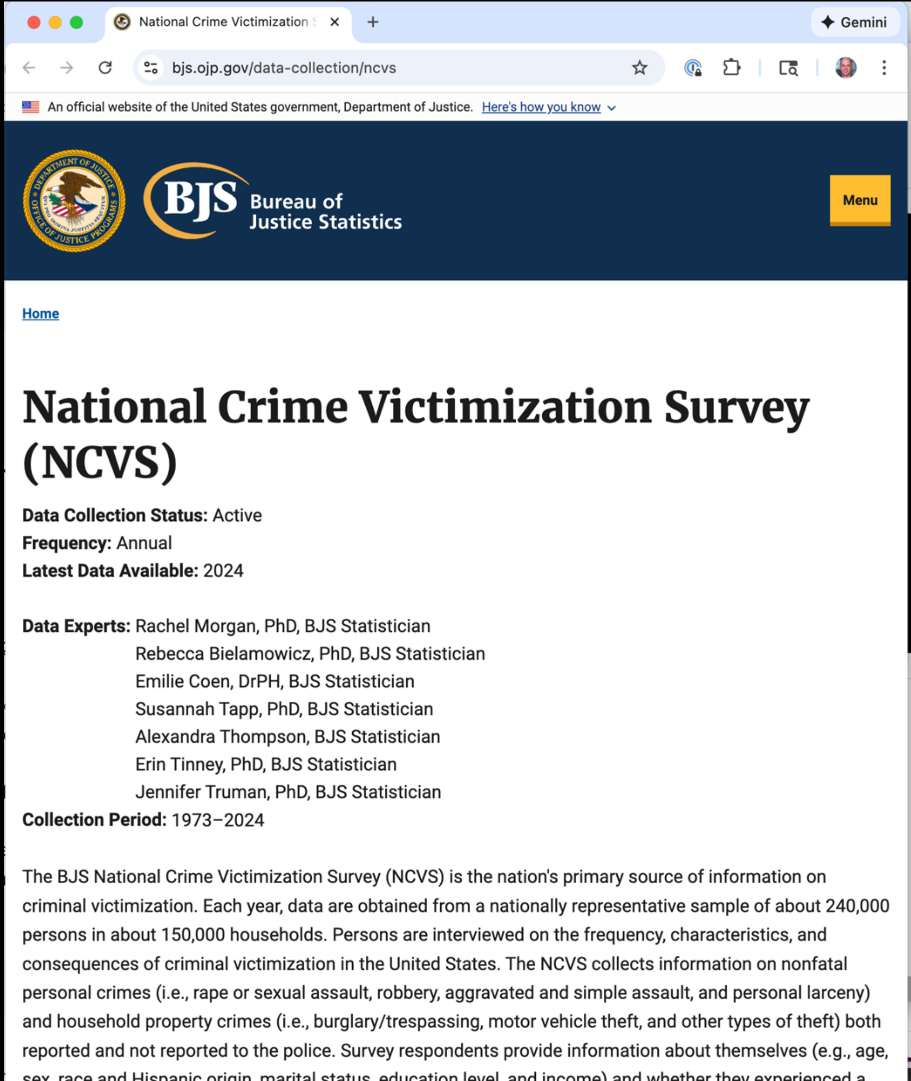
National Crime Victimization
Gemini
bjs.ojp.gov/data-collection/ncvs
An official website of the United States government, Department of Justice. Here's how you know
BJS Bureau of Justice Statistics
Menu
Home

National Crime Victimization Survey (NCVS)

Data Collection Status: Active
Frequency: Annual
Latest Data Available: 2024

Data Experts: Rachel Morgan, PhD, BJS Statistician
Rebecca Bielamowicz, PhD, BJS Statistician
Emilie Coen, DrPH, BJS Statistician
Susannah Tapp, PhD, BJS Statistician
Alexandra Thompson, BJS Statistician
Erin Tinney, PhD, BJS Statistician
Jennifer Truman, PhD, BJS Statistician
Collection Period: 1973–2024

The BJS National Crime Victimization Survey (NCVS) is the nation's primary source of information on criminal victimization. Each year, data are obtained from a nationally representative sample of about 240,000 persons in about 150,000 households. Persons are interviewed on the frequency, characteristics, and consequences of criminal victimization in the United States. The NCVS collects information on nonfatal personal crimes (i.e., rape or sexual assault, robbery, aggravated and simple assault, and personal larceny) and household property crimes (i.e., burglary/trespassing, motor vehicle theft, and other types of theft) both reported and not reported to the police. Survey respondents provide information about themselves (e.g., age, sex, race and Hispanic origin, marital status, education level, and income) and whether they experienced a

Source:

National Crime Victimization Survey Identity Theft Supplement

- Followup questionnaire to NCVS
- Done roughly every two years
    - Last done in 2021

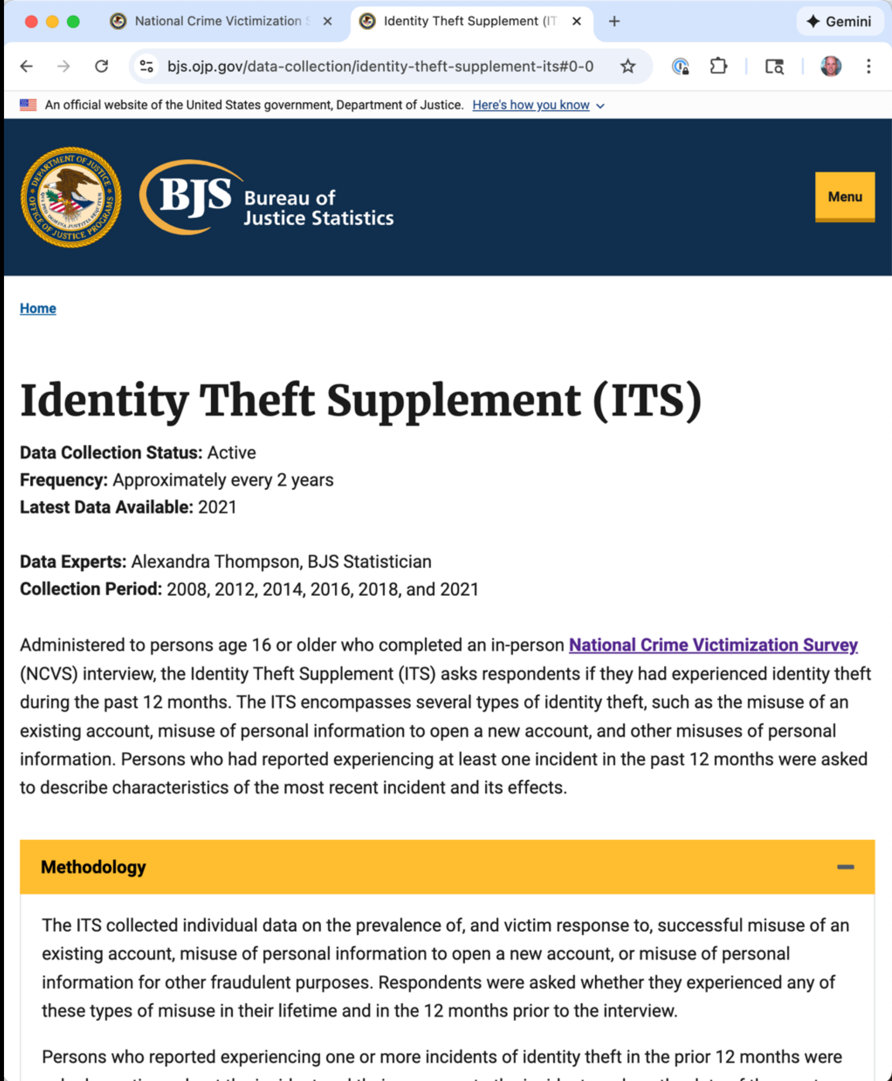
National Crime Victimization
Identity Theft Supplement (IT
Gemini
bjs.ojp.gov/data-collection/identity-theft-supplement-its#0-0
An official website of the United States government, Department of Justice. Here's how you know
BJS Bureau of Justice Statistics
Menu
Home

Identity Theft Supplement (ITS)

Data Collection Status: Active
Frequency: Approximately every 2 years
Latest Data Available: 2021

Data Experts: Alexandra Thompson, BJS Statistician
Collection Period: 2008, 2012, 2014, 2016, 2018, and 2021

Administered to persons age 16 or older who completed an in-person National Crime Victimization Survey (NCVS) interview, the Identity Theft Supplement (ITS) asks respondents if they had experienced identity theft during the past 12 months. The ITS encompasses several types of identity theft, such as the misuse of an existing account, misuse of personal information to open a new account, and other misuses of personal information. Persons who had reported experiencing at least one incident in the past 12 months were asked to describe characteristics of the most recent incident and its effects.

Methodology

The ITS collected individual data on the prevalence of, and victim response to, successful misuse of an existing account, misuse of personal information to open a new account, or misuse of personal information for other fraudulent purposes. Respondents were asked whether they experienced any of these types of misuse in their lifetime and in the 12 months prior to the interview.

Persons who reported experiencing one or more incidents of identity theft in the prior 12 months were

Victims of Identity Theft, 2021

Lower prevalence (but not the same methodology, questions)

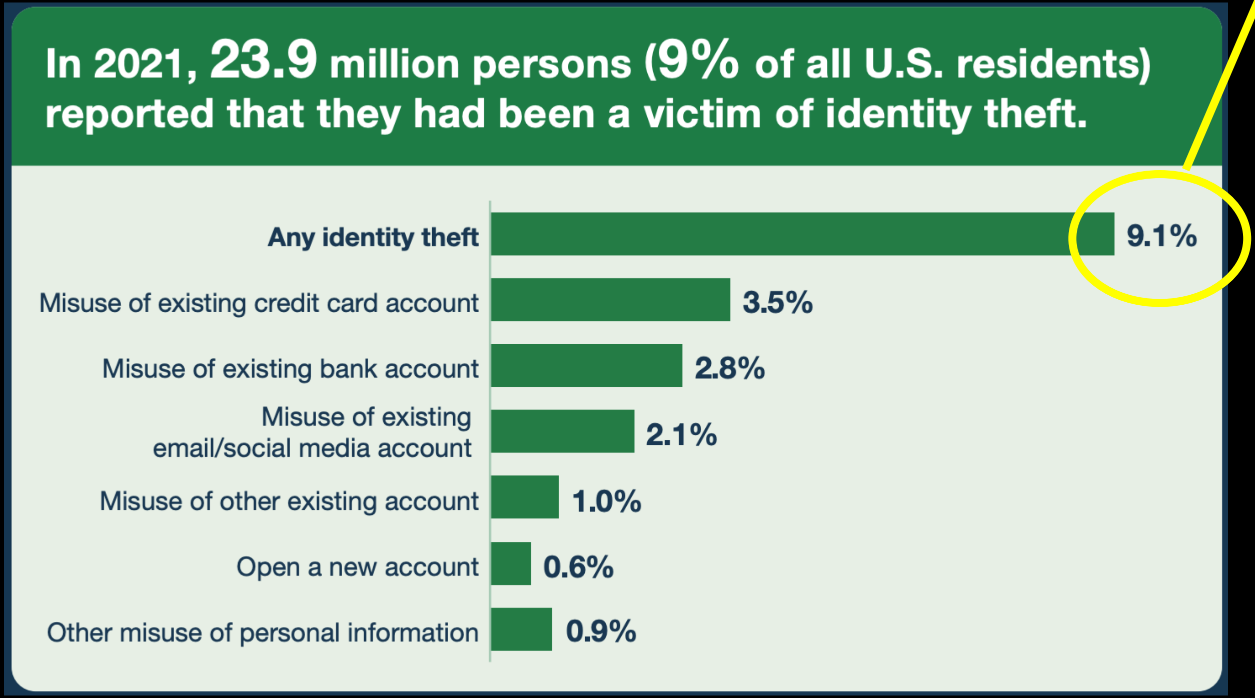

In 2021, 23.9 million persons (9% of all U.S. residents) reported that they had been a victim of identity theft.

Comparing other data sources

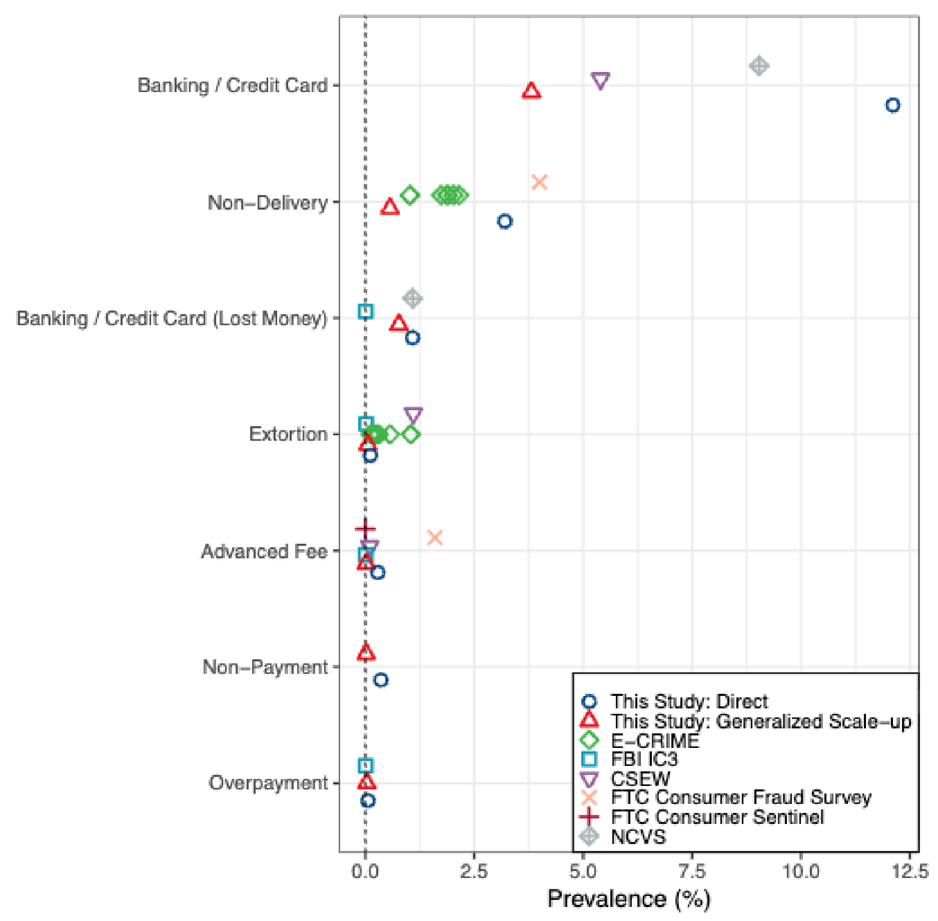

RQ2: Monetary Impact

Median losses are modest

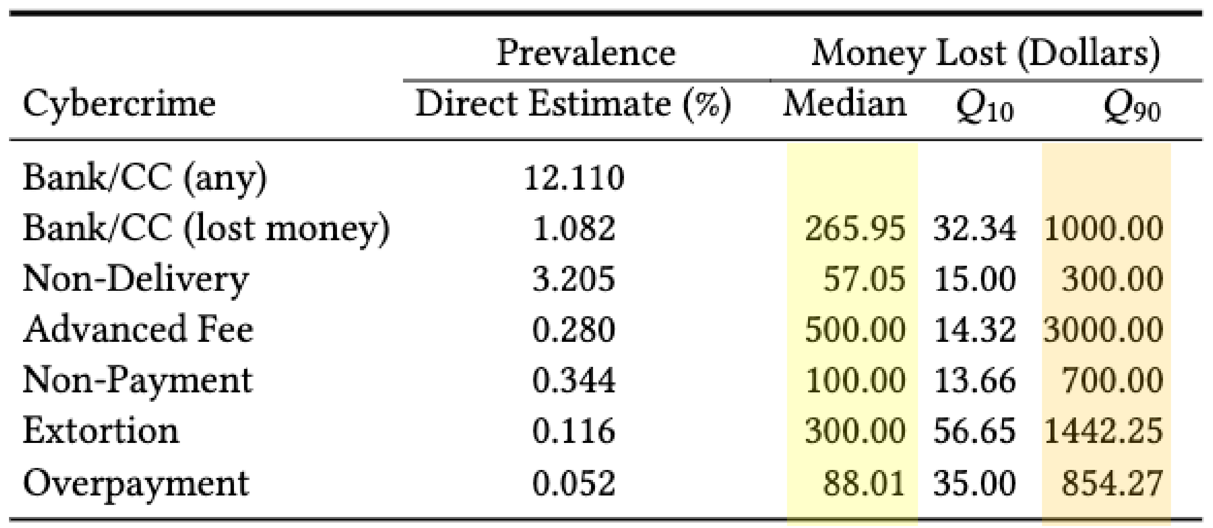

	Prevalence	Money Lost (Dollars)		
Cybercrime	Direct Estimate (%)	Median	Q_{10}	Q_{90}
Bank/CC (any)	12.110			
Bank/CC (lost money)	1.082	265.95	32.34	1000.00
Non-Delivery	3.205	57.05	15.00	300.00
Advanced Fee	0.280	500.00	14.32	3000.00
Non-Payment	0.344	100.00	13.66	700.00
Extortion	0.116	300.00	56.65	1442.25
Overpayment	0.052	88.01	35.00	854.27

Distribution is heavily right-skewed

RQ3: Demographic Inequities

Look ma! No LASSO!

Weighted logistic regression with demographic predictors: age, race/ethnicity, gender, income, education, internet skills

Key findings:

- **No single “victim profile”** across all crime types — patterns vary
- **Internet skills** matter: higher skill associated with different victimization patterns
- Some crimes show income/education associations, but direction isn’t always intuitive

All associations are **observational**, not causal!

Limitations

- **COVID-19 timing** (July–September 2020) may have affected both cybercrime rates and survey participation
- **English-only** survey misses non-English-speaking populations
- **Recall and recognition bias** — victims may not accurately recall or identify experiences
- **Cross-sectional** — single snapshot, no ability to measure trends
- **NSUM** (not covered) **underperformed** — barrier effects made it produce underestimates rather than the hoped-for improvement

Some Takeaways

1. **Cybercrime on individuals had modest effect** — Most crimes are rare, but some are common (credit card fraud), and in general cost users little

   “a typical consumer stands to lose $22.90 (= $6.87 ÷0.3) and in the worst 90th percentile case a consumer stands to lose $111.87 annually”

2. **Corporate estimates of crime seem much higher** — Gartner, McAfee, Norton, …

Backup: Network Scale-Up Method (NSUM)

NSUM: What It Is & Why You’d Want It

Problem: Direct surveys only capture victimization the respondent **recognizes and is willing to report**.

NSUM idea: Ask people about **people they know** instead of themselves:

> “How many people do you know who experienced [cybercrime X]?”

Why this might help:

- Even if *you* wouldn’t report your own victimization, friends/family might mention theirs
- Aggregating across many respondents’ networks could capture hidden victimization
- Has worked in **public health** for estimating hard-to-count populations (IV drug users, sex workers)

How NSUM Works

1. **Estimate each respondent's network size** — "How many people do you know named Emily? Walter? Ralph?" (12 calibration names with known population frequencies)
2. **Ask about the target** — "How many people do you know who experienced [crime X]?"
3. **Basic estimate:** Prevalence ≈ (reported crime contacts) / (total network size)
4. **Correction factors** (Generalized NSUM):
   - **Degree ratio:** Do victims have larger/smaller networks?
   - **Visibility factor:** What fraction of a victim's network *knows* about the crime?

What We Expected vs. What Happened

Expectation: NSUM estimates should be **higher** than direct self-reports (capturing hidden victims, like in public health).

Reality: NSUM estimates were **lower** in every crime category.

Crime Type	Direct Estimate	NSUM Estimate
Banking / Credit Card	~12.1%	~3.9%
Non-Delivery	~3.2%	~0.6%
Non-Payment	~0.3%	~0.02%
Extortion	~0.1%	~0.06%

What NSUM’s Failure Tells Us

NSUM is not a universal fix for under-reporting. It requires that network members *know about* each other’s status.

Why it works in public health: Being an IV drug user or sex worker is often known/observable to close contacts.

Why it fails for cybercrime: Victimization is private, brief, embarrassing — invisible even to close contacts.

The silver lining: This is itself a finding — cybercrime is a largely *private* experience, which means:

- Social support for victims may be lacking
- Community-based interventions face adoption challenges
- Future NSUM attempts could focus on closer contacts (family, close friends)