INAUGURAL WORKSHOP ON CYBER PUBLIC HEALTH

Workshop Report [24-01]

A report that captures the discussions of future directions and research in the nascent field that applies public health concepts to the area of cybersecurity.

CyberGreen Institute
June 2024
DOI: 10.13140/RG.2.2.22399.62887

EXECUTIVE SUMMARY

The first workshop on Cyber Public Health was an exciting event held in New York in January 2024. Participants engaged with a new way of thinking about cybersecurity, focused on the idea that we should measure the health of populations, rather than focus on the more traditional perspective of specific attackers versus defenders. We explored how this perspective leads to new ways of measuring harms, who's impacted by those harms, and data sources we might tap into or create to better understand these harms. Some of the interesting ideas include:

- Cybersecurity failures can create harm to "Digital Activities of Daily Living[1]" where people can't do what they want or need, either because they're blocked by proactive defenses or because they have to clean up after a compromise of a system ("malware/ransomware") or account ("account takeover").

- We heard a perspective that health is about mindless adversaries and security is about intelligent ones. However, health is not just about pathogens; people can be hurt by actions (their own or other people's) and systems we create or operate. For example, smoking tobacco damages the health of both the smoker and those around them. So, the perspective "health is about mindless adversaries and security is about intelligent ones" is less true than we might have expected.

- Moreover, smoking is promoted by companies whose owners are protected from liability despite the harms they cause. The companies that promote and sell spyware causing harm (often to specific populations such as journalists or activists) are similarly protected from liability. In public health research, we might seek to gather data on how many people are actively (and passively) exposed to tobacco smoke. How would we go about gathering such data on the prevalence of spyware on journalists' phones?

This report is designed to serve three audiences. First and foremost are the participants in the workshop, to provide them with a record and path forward. Second are other scientists, including specialists in public health and computer science, who may find these ideas interesting or provocative. The third is policy specialists, who may find the Cyber Public Health metaphor a useful frame for policy, both measuring the efficacy of security and interventions.

This report starts with a brief explanation of the workshop and some background on the concept of Cyber Public Health. Section 2 captures workshop output and notes: the conversations we had, organized into a useful record, and captures research questions. Section 3 explores directions we

[1] Adapted from the medical concept "Activities of Daily Living" (Edemekong, Peter F., et al.)

may take, while section 4 identifies some of the challenges. Section 5 provides a list of the participants and finally, section 6 provides a bibliography.

1. THE WORKSHOP: WHAT WE DID

On January 9, 2024, CyberGreen and Google co-hosted the inaugural Cyber Public Health workshop as an exploration into future directions and research in the nascent field that applies public health concepts to the area of cybersecurity. The hybrid workshop was an opportunity to assess past achievements, introduce fresh perspectives to the discourse, and establish shared objectives for collaborative research and development.

Approximately 35 participants attended the event from various backgrounds, including cloud providers, academia, government organizations, professionals in public health and cybersecurity, and representatives from security vendors and community organizations.

We had two short keynotes on public health (Nathan Taback, University of Toronto) and Cyber Public Health (Adam Shostack, University of Washington and CyberGreen). These were followed by two breakout working sessions, one on definitions and one on data, each followed by a readout which allowed us to share insights from the breakout sessions to the larger group. Finally, we held an organizing session in which we asked each person for next steps, then grouped them.

The workshop looked at units before harms; a participant pointed out that had we done it as harms then units, we would have learned more. While we can't adjust what we did, we note the lesson, and present the workshop's output in the more useful order. These two sessions produced output that we now break into three groups:

- Harms
- Definitions
- Data sources

We facilitated conversation by providing some guiding questions to breakout groups. We used the Chatham House Rule, so the report should be read as a record of the event and points raised, not as a consensus of the participants.

1.1 BACKGROUND: CYBER PUBLIC HEALTH AS A NEW APPROACH

As can be deduced from the epidemic of computer breaches, compromises, and disruptions, there are persistent problems in cybersecurity. In exploring solutions to those problems, perhaps new perspectives can be helpful. Preliminary exploration of applying public health concepts to cybersecurity seems promising.

It is difficult today to characterize the health status of the internet and its connected services, devices, and people. There are no clear and universal definitions for describing the entities we aim to

keep healthy – whether they are people, computers, or accounts. There is no consensus on what constitutes their morbidity or mortality.

On an individual service, device, or account, we can measure health (e.g. service is operating as expected, malware is not present, no unauthorized account access). There may even be statistics that "turning off automatic updates increases the incidence of new malware by 'X'," but such statistics are complicated. For example, Microsoft delivers their Malicious Software Removal Tool each month via Windows Update. So, if updates are off, new malware may not be detected.

There is no way to measure what specifically confers health benefit, so we cannot tell what interventions lead to good outcomes and avoid bad outcomes. We have no methods to make public internet health decisions – like when to isolate or constrain a small population to preserve the safety of a broader population (e.g. the equivalent of the mandatory quarantine of antibiotic-resistant, TB-positive humans).

For example, recently, the American Hospital Association advised member hospitals to disconnect from Change Healthcare, a payments processor and a victim of a ransomware attack. They presented no evidence about how or why it would be protective. However, there has been a very clear impact on both care delivery and payments as a result.

Cybersecurity models in use today often use analogies that focus on the idea of an adversary in crime, espionage, or war. They fail to look at both the larger context in which the "battle" is unfolding (i.e. are there others also fighting similar adversaries, or is there some collective action we could take that would help us all?), nor do they help guide preventive actions and behaviors which, when taken, bring benefit to the larger population.

In Cyber Public Health, where we apply public health models and tools to cybersecurity, data is critical. While there is much to be gained by scanning public resources, significant online activity remains under private control. Moreover, many compromises and system failures are not disclosed and so little can be learned collectively about what created a particular vulnerability, how it was exploited, what provided a cure, and what can ensure future prevention. To move this field forward, we need the engagement of commercial, government, and non-profit organizations to help both deepen and broaden datasets to allow us to understand and track internet health risk factors.

Nearly 30 years ago, the world set out to understand the human genome – to help understand the composition of what makes us human and what causes disease. It was a significant computational problem, but one that was able to be solved with a global collaborative effort. The time to make such an effort for global Cyber Public Health is here. With collaboration, data collection, and applying new ways of thinking about cybersecurity, we can build a more resilient internet.

1.2 THE METAPHOR AND THE CONCRETE

Medical metaphors have been in use in computer security since at least 1984 when Fred Cohen wrote about "computer viruses." Readers may have their own favorite examples, all of which may beg the question: does another metaphor help?

THE MEDICAL AND PUBLIC HEALTH METAPHORS

These metaphors help to make the complex and often opaque aspects of cybersecurity explainable to a broad spectrum of society. The cybersecurity as health metaphor can also be empowering as it encourages organizations as well as individual users to maintain basic standards of good "cyber hygiene". For example, by updating and patching software regularly, and adopting good habits in the use of passwords. This metaphor has appeared widely across guidance issued by the UK National Cyber Security Centre (NCSC) and is embedded in its Cyber Essentials scheme. Microsoft also runs a health dashboard and uses the health metaphor widely across its services. The research base supports these approaches: "the health metaphor suggests that preventative measures may be much easier and cheaper to implement than after-the-fact care […] Like diseases, computer security threats can spread rapidly, evolve […], and be addressed with defensive measures ranging from preventative care to treatment and quarantine of active infections" (Wolff).

Yet, the "cyber hygiene" metaphor risks placing the burden of responsibility – and even the blame – upon individuals when things go wrong (Slupska). According to one study: "As in discourses of epidemics and contagion, cyber insecurities are generated by individuals who behave irresponsibly thus compromising the health of the whole" (Hansen and Nissenbaum). Research into the folk models that people display in their attitudes to cybersecurity reveals that the health metaphor can also imply passivity and inevitability: cyber insecurities and viruses are simply something that are all around us and that we can "catch" at any time – like the common cold (Wash).

Despite these concerns, there remains broad enthusiasm for the public health metaphor, including proposals that devices could be required to present vaccine or "health certificates":

> To improve the security of the Internet, governments and industry could […] engage in more methodical and systematic activities to improve and maintain the health of the population of devices in the computing ecosystem by promoting preventative measures, detecting infected devices, notifying affected users, enabling those users to treat devices that are infected with malware, and taking additional action to ensure that infected computers do not put other systems at risk. (Charney)

However, it is important to remember the limitations of "infection" within this public health metaphor. Biological viruses and pathogens are indiscriminate and adapt and evolve slowly

(following Darwinian processes) but computer viruses and their vectors are targeted and designed by human adversaries to adapt fast (following Larmarkian processes).

The CyberGreen Institute is a longstanding champion of empowering people and organizations to take proactive measures to help them avoid and mitigate cybersecurity issues, rather than focusing reactively upon treating threats and responding to attacks: "Such approaches are analogous to treating a case of malaria through medicine, while leaving the nearby mosquito swamp untouched or developing cancer treatment technology while paying little attention to the population's tobacco use" (CyberGreen). CyberGreen also takes inspiration from the metaphor and from the public health sector by advocating for better data gathering, tracking, and publishing of cybersecurity metrics – in the same way that public health bodies use population health studies (CyberGreen).

THE CONCRETE

This workshop was focused on advancing past these metaphors. We aim to bring together a community who will take these ideas and explore ways to implement them.

Public health has a toolbox, including epidemiology, behavioral sciences, data gathering and analysis infrastructure, and public policy, all of which may be concretely instructive. For example, exploring the concept of "prevalence" (the fraction of a population with a condition) in Cyber Public Health allows us to take note of several things, including the lack of population data and the lack of definition of "malware families." These observations may not seem inherently valuable, but they open a door for further consideration: is malware getting worse? We have generally studied that question by examining what the malware does, rather than the "incidence" (the fraction of a population that acquires a condition in a given time) of new malware. Cyber security has tended to avoid that because its measurement is difficult, but in 2011, the lead author was able to use incidence data combined with propagation information to drive Microsoft to ship a fix for Autorun in Windows Update, leading to a dramatic drop in malware infections and cleans for over a year (Microsoft, "Zeroing in").

In partial contrast to computer science, public health might be seen as a 'socio-technical' science, whose field of study includes the humans whose health they seek to improve. This leads to a field of practice where the success of an intervention is a primary metric: does it help people live longer, healthier lives?

Vital Signs and Vital Statistics

We ran into confusion over the difference between vital signs and vital statistics. When we go to the doctor, they measure vital *signs*: body temperature, pulse rate, blood pressure and breathing rate. Each of these can be measured quickly and easily and together, they give useful insight into acute problems and treatment effectiveness. Vital signs and vital statistics provide valuable information, but they serve different purposes. Vital signs are immediate, direct measurements of basic body functions and are used primarily for medical assessment and monitoring. On the other hand, vital

statistics are collected and analyzed over time to study population trends and inform public health initiatives. They do not directly relate to each other, but both contribute to the overall understanding of individual and population health.

We will continue to identify and adapt specific tools from the field of public health, while also leveraging its metaphors and framing. We feel this presents a compelling alternative to the traditional military or espionage framing often used for cybersecurity.

1.3 KEYNOTE SUMMARIES

Dr. Nathan Taback presented "A Data Science Perspective on Cyber." He started from the idea that public health is a population-oriented discipline, and contrasted it with medicine, which focuses on the course of care for an individual patient. He discussed a set of core sciences, including prevention effectiveness, epidemiology, lab science, informatics/data science, and surveillance. He highlighted the importance of data governance, how data is defined, collected, and normalized, along with stewardship and access issues. Then he raised the question "what problems are we trying to solve?" and presented a concept of "epicycles of analysis."[2] We can consider both prospective and retrospective analyses. We can look for direct outcomes, which are often hard to measure directly, or dependent variables, which are often easier, but carry a risk of measuring the wrong thing, measuring something which is confounded, et cetera. To illustrate his point about dependent variables, he gave the example of measuring patient survival after cancer treatment. We hope that takes a long time, and so instead of measuring survival, tumor size is measured as a surrogate. This is convenient, but over time some might forget that tumor size is correlated with survival, but ultimately measures something different. He discussed the well-reported story of Google Flu Trends and discussed its flaws, well known to statisticians, data scientists, and epidemiologists. An analysis of 50 million queries ultimately yielded a dataset of 1,100 doctor visits. In retrospect, Google Flu Trends was found to have predicted twice the number of doctor visits observed. Dr. Taback's presentation is available.

Adam Shostack presented "Towards a Science of Cyber Public Health." The talk started by wryly noting that no one is saying "Cybersecurity is going great, keep it up," despite large investments. He pointed out that we don't know if we have more incidents than last year, and asked: if we do, are we seeing more growth in incidents or population? He told a story about fixing Autorun.[3] From there, he transitioned to a brief description of a public health toolkit of measuring important data, discovering harms to people, and investigating clusters to fix the causes via policy and behavioral change. He closed by saying that public health can complement other frames like cyberwar and cybercrime, and talked about going beyond the metaphor, drawing on the story of the Cyber Safety Review Board. The concept of "an NTSB for security" had been around for 30 years. Work done

[2] See (Peng and Matsui) for an example.
[3] See (Microsoft, "Zeroing in").

with Rob Knake to organize the ideas, objections and tradeoffs led to the establishment of a real NTSB for security, and we can do the same for Cyber Public Health. Mr. Shostack's presentation is available.

1.4 WORKSHOP OUTCOMES

The inaugural Cyber Public Health workshop was a thought-provoking event. We achieved our primary objective: to articulate the promise and challenges of applying public health methods and frameworks to cybersecurity. The workshop brought together a multidisciplinary group. We generated substantive insights and meaningful connections and advanced the conceptual and practical foundations for the emerging discipline of Cyber Public Health. Specific outcomes included:

- Establishing a shared understanding of the core concepts and potential applications of a public health approach to cybersecurity.
- Delineating the types of harms, data sources, and units of analysis most relevant to assessing and improving population-level cybersecurity.
- Identifying critical research questions and data needs to advance the field.
- Building relationships to sustain cross-disciplinary dialogue and collaboration beyond the workshop itself.

This report serves as an actionable record of the workshop and presents research questions that we hope will be tackled over the coming years. Doing so will require further collaboration, funding, and hard work.

2. WORKSHOP OUTPUT AND NOTES

2.1 HARMS

There are many harms which arise from a lack of cybersecurity. They include **direct harms**[4]: people unable to access services, loss of access to sensitive data and intellectual property, businesses unable to serve customers, governments unable to serve citizens or visitors. When these harms are a result of denial of service, no one other than the attacker profits. The costs here are greater than the gain to attackers.

In contrast, other harms are **unauthorized** or **unearned** transfers as a result of theft or scams. Here, the financial cost to the organization is roughly equal to the gain to the attacker, excluding the cost of cleanup, and excluding the cost of privacy for the organization's customers. Yet other harms

[4] A participant suggested that the term "direct harms" is not quite right and that these are harms to availability, integrity or confidentiality.

are transfers made **under threat** such as ransomware or blackmail. Blackmail may be less obvious, but often[5] impacts those who have sent naked pictures of themselves. While the sexting example is salient, blackmail is a threat to someone's privacy, and is included in Solove's taxonomy.

There are also harms which arise when people don't engage fully in things they want to do, or when they engage in security toil rather than other activities. If we use the availability frame, these precautionary or avoidance harms are also harms to availability.

Lastly, there are costs when time or money is spent on security rather than something of more direct value. Participants considered a wide range of harms, which are captured below.

HARMS SUGGESTED BY GROUPS DURING THE WORKSHOP

- Impact to "Digital Activities of Daily Living" were put forth as a category of harms. Another group chose to list some different types of harms, such as technical harms, breaches, violations of CIA, harms to trust, risk conditions, etc.
- Impact to intimacy. We discuss blackmail above, and we can suspend judgment or prurience and consider that if people want to engage in sexting but don't because of security concerns, that is an instance of them not fully engaging in things they want to do.
- Impact to liberty. People may not engage in political discourse, may not report a crime, or other activities because of fears of surveillance, breaches, or other, less clearly expressed concerns.
- Public health has a concept of "Global burden of disease"[6].
- Disability Adjusted Life Years" (DALY) is a measure of the burden from chronic disease. The number of years living with a condition is counted and adjusted with a multiplier of 0 to 1. A multiplier of 0 indicates no impact, while a 1 is as bad as dying. DALY measurements in public health are in units of years, not the economic value of those years, or the cost of care. There are economists who measure those costs.
- Even in cyberspace, should there be a focus on harm in the physical world? (e.g. power goes out, people can't eat, can't get money, or perform basic functions…)
    - How do we count accurately in such cases? (e.g. a power outage impacts all of my accounts at once)
    - How do we engage with substitutability? For example, a person being able to "get money" (perhaps from a bank or bank machine) might be substitutable by another bank account, by use of a credit card, or a personal loan.
- Public health has a concept of diseases of poverty and diseases of affluence. Perhaps there are equivalents in Cyber Public Health?
- There were questions pertaining to legal issues. We explored impact in the context of harms.

[5] We lack statistics on its frequency, though it is widely reported that "sexting" is a normal part of modern dating. Thus, we both lack data on incidents and agreement at the population level: do we estimate those who sext? Single people?
[6] See https://www.healthdata.org/research-analysis/gbd and https://vizhub.healthdata.org/gbd-compare

    - Insurance companies quantify loss only in financial terms, and at some point, all harms - even emotional distress, grief - can become financial harms.
    - Death can be treated as financial harm, and public policy has various valuations of a human life.[7]
- Issues with current systems were brought up, with one example related to Common Vulnerability Scoring System (CVSS) and the breadth of its coverage.
    - CVSS covers the technical depth of the impact of a vulnerability, but doesn't cover the breadth of how this vulnerability could impact a sector, a state or nation.
- Vulnerability data is collected by a variety of actors including Shodan, Shadowserver and CyberGreen. Is there a way to correlate problems to vulnerability?
- There was extensive discussion of the need to refine definitions, collect data, and establish the institutions to do so.
- One group discussed a need to understand and articulate the difference between harms and risks.
- One group listed some different types of harms:
    - Technical compromises and the organizational/data consequences associated
    - Breaches (account takeover)
    - CIA (confidentiality, integrity, and availability) triad compromise
    - Means of harm (e.g. successful phishing compromise but no attack)
    - Actual technical/physical damage
    - Risk conditions (getting credentials, having knowledge of a vulnerability)
    - Trust of or within an organization
- A group listed types of data, along with who might have it:
    - incident/crime data (enterprises, investigators, regulators, cloud providers, commercial/non-commercial collectors)
    - Vulnerability data (enterprises, platform and cloud providers, non-commercial scanning entities)
    - Malicious traffic data (telco and transport, cloud providers, DNS entities, honeyfarms)
- Some questions worthy of further consideration:
    - If a machine is infected but has no negative impact, does it matter?
    - Is it a form of harm where the villain benefits but the victim doesn't care?
    - What is the cyber equivalent of a person killed?
    - How do we define a proportionate response?
    - What could we learn from VirusTotal, the largest malware database in the world, relying on users to submit samples of malware, anonymously.

[7] There are surveys at http://www.law.harvard.edu/programs/olin_center/papers/pdf/Viscusi_517.pdf or https://en.wikipedia.org/wiki/Value_of_life

- One public health equivalent was a project where individuals submitted their Explanation of Benefits and it showed how your benefits could apply differently depending on the provider you saw; this led to price transparency.
    - This led to a more concrete discussion on cost: There are financial costs associated with logs - both storage and analysis - which could be prohibitive, depending on the entity.
- We might learn about how important formally-aligned naming conventions for malware are; perhaps most samples are categorized into the same names by most engines?

Some harms may touch on sensitive areas. In public health, there are highly sensitive harms (such as sexually transmitted disease) which are reported within carefully designed regimes. There are other harms (such as motorcycle injuries) which are unregulated or regulated unevenly across states and countries. Companies hate being forced to report breaches, and the list of reporting requirements is growing in the United States. In addition to personal data breaches, the SEC requires reporting by public companies, the DoD requires reporting from the "Defense Industrial Base," and CISA has proposed reporting by a broadly defined set of critical infrastructure companies and their suppliers.

Research questions include: What are the sensitivities? Can we leverage these reporting requirements, crime reports, or other data sources? What is the impact of sector-by-sector and state-by-state regulation? Does it create opportunities for experimentation or hinder them by imposing complexity?

Organizations can believe it's advantageous to ensure they don't know about incidents or breaches, since knowing but not taking action is more harmful to the organization's reputation than not having logs enabled (not knowing). It may be that there is a policy argument for limiting the liability that comes with knowing about an incident to improve public health. Other economic arguments were put forward:

Cyber offsets (like those we have with energy efficiency) because we believe it is a public good and it is therefore incentivized with tax credits and a liability shield.

- Service providers could also do it on an organization's behalf in the same way that HR Block does your taxes for "free" but knows a certain percentage of the population will get a return and HR Block will keep a percentage of that return.
- ISPs/CSPs/other service providers (e.g. Google for GMail) should have an incentive to report out information. We already see something similar with credit card fraud reporting or Google's Transparency Reports or Exposure Notifications for COVID.

The idea of not knowing seems easier in cybersecurity, since we lack dead bodies. Public health specialists are able to measure "excess deaths" and use that to infer unmeasured deaths from the pandemic. The technique is not free of controversy.

POSSIBLE HARM-ORIENTED RESEARCH QUESTIONS:

- Is there a standard list of harms used by public health professionals, or a list of lists?
    - We believe there is a mix of commonly agreed items, such as disease, emergent issues (food deserts) and situationally useful ones (accidents which lead to death in the elderly).
- In cybersecurity and with privacy, there are unique harms.
    - **Security harms** tend to center around confidentiality, integrity, and availability (the CIA triad), and
    - **Privacy harms** tend to center around data subject rights and violations of trust.
    - But there is another realm of harms that really arise from the nexus of both: when there are both security and privacy harms which happen at a population level. These are **societal harms**: when there is a loss of trust by populations of users, where there is social unrest, repression and attacks on expression, economic disruption, cultural isolation and distrust. These societal level harms are important to understand as they undermine the digital lives of societies.
- With sufficient population-based data, can the likelihood of harms be predicted?
    - One aspect of public health data collection is that it can permit predictive modeling of outbreak behaviors. In the context of cybersecurity, can similar predictive modeling be developed to allow for preventative action?
- For Cyber Public Health, we have organized workshop output into the harms above.
    - Is the list generally useful as a starting point for other research?
    - What are its strengths, limits or weaknesses?
    - Are there scenarios where it either shines or shows weakness?
- Assessing the impact of repairable harms (e.g. stolen credit card number) and irreparable harms (e.g. mental health, death, social manipulation).
    - How do victims of credit card fraud recover financially and emotionally compared to victims of severe mental distress or social manipulation?
- Comparative analysis of immediate harms and delayed harms. In cases of social manipulation, what are the immediate observable effects versus the long-term societal impacts?
- Legal and regulatory responses to different harms: what are the differences in legal recourse available for repairable harms versus irreparable harms?
- How should we be categorizing sensitivities around either harms or units?
- Can we leverage reporting requirements that center other goals to gather data on harms? Which regimes help most or least and in what ways? What do the helpful reports tell us?

What changes might be required? How hard are those changes? What could we achieve in collaboration today? In a few years? What would require legislative changes?

2.2 DEFINITIONS

Our definitions now cover a mix of atomic units and the relationships between them, the harms which may affect them, and the types of units they are. These are working definitions. They may well overlap, contradict each other, be more or less useful in various situations, etc. A key research activity will be to discover which ones are most broadly applicable (like prevalence and incidence), which ones are useful in part due to imprecision (like health belief models), and which ones can be implemented in software (versus, say, surveys).

ATOMIC UNITS

In public health, the fundamental unit is people, considered as a population. The equivalent unit for Cyber Public Health is not obvious. Perhaps it's a person, but a person doesn't get sick in cyberspace. One or more of their accounts do, and those accounts are frequently protected differently, making them an interesting atomic unit.

Perhaps it's a computer? If so, do we care? A computer, in theory, can be re-booted, re-initialized or have its operating system re-installed and brought back to a 'clean state.' A computer is not a unique being with a personality (or soul) worthy of preservation or endowed by its creator with any inalienable rights. Therefore, why would we care about its health beyond the risk/benefit it provides to people? We care because of the impact it may have on people if it is "diseased."

In preliminary work (Shostack), we identified computers (including IoT devices and smartphones) and accounts as primary units, and in the workshop, we re-opened the question. In drafting this report, we explore the term "atomic unit." The English word *atom* derives from the Greek for indivisible or uncuttable. The proposal references that, and the notion that they can be combined.

In public health, core vital statistics are births and deaths, both carefully tracked. Population is derived from those (with adjustments for immigration), and cross-checked against census data. In Cyber Public Health, creation data (analogous to births) may be more available than death data. There are complex challenges in tracking deaths, including a final power-down or disposal of a device, the closing of an account, the deletion of a VM from a VM store.

This brings us to a fundamental research question: What constitutes the vital statistics of Cyber Public Health and what challenges exist in gathering each?

Simple units:

- Technical units

    - Devices
    - Virtual Machines
    - Malware(?)
- Human units
    - Accounts
    - Users

TECHNICAL UNITS

Devices have the convenient properties of being manufactured and sold, which are analogous to birth. But some devices are never sold, others are never used or used only briefly. Other data may be available telemetry data on devices connected to the internet (e.g. phone companies know the number of phones on their network). There is also a great deal of automated endpoint device reporting: login attempts, breach attempts for emails, or anything indicating something less than the desired state.

Malware may also be measurable as "atomic units", similar to the way diseases are measured. Diseases tend to be measured by consensus techniques, ranging from use of the DSM, a handbook for the diagnosis of mental health disorders, to lab tests (which encompasses a wide-ranging set of techniques). In contrast, the identity of a malicious program is defined by trade-secret rules maintained by individual vendors, and the vendors have not contributed to standardization efforts. For example, the Common Malware Enumeration project went dormant in 2007.

Other technical programs may produce interesting units. For example, standards such as the Open Cybersecurity Schema Framework (AWS, Splunk, and other vendors) that help normalize and standardize log formats.

HUMAN UNITS

An atomic unit *could* be the human, but humans are not digital, and there are ways in which it's easier to measure our digital representations. Maybe the equivalent is accounts (measuring the number of compromised accounts, user growth numbers in financial reporting, etc.) On the other hand, one unhealthy account doesn't necessarily impair you because you can use other accounts/devices. A person can be analyzed through the lens of their account with a provider, and most people have a relationship with several major entities like Google and Facebook. A person may well have several accounts with each of these.[8] When these are usefully counted as separate or "de-conflicted" is situational. There may be times when Digital Activities of Daily Living are impacted because someone can't log in to a specific account, and other times when life goes on.

[8] For example, author Shostack has at least a personal Google account, one for CyberGreen, whose domain is hosted by Google, and two via the University of Washington, which has a Google domain for the University and one for the Computer Science department.

Is a person in cyberspace the sum of their accounts? Are some accounts more critical than others? Are they common across people or populations? There are debates going on right now about the harm of social media accounts on teens -- so the presence of or access to a service is seen as potentially harmful to health.

Is an account like a cell in a body? Compromise of a single cell rarely results in disease but, depending on circumstances, can be fatal (i.e. a cell becoming cancerous). Similarly, the compromise of a single account may or may not be troublesome, but can lead (if "untreated") to other problems. This implies people in the context of cyberspace are the aggregate of their "accounts"[9].

We should not ignore that these are representations; techniques including surveys, interviews, and incident or crime reports give us visibility into human views.

COMPOUND UNITS

- Software (open source, components)
- Applications
- Services
- Networks
- Enterprises, including
    - The individuals who work there, their offices and homes
    - Their stakeholders including
    - Their customers and the data the enterprise holds on those customers
- Third parties
- Digital identity
- Governments
- Attackers

Software and applications are listed separately as software might be the xz library, or sshd which sometimes incorporates it as a result of a packager, such as RedHat, integrating it with systemd. An example of an application might be Quickbooks (either classic or online); both are likely the result of Intuit acting as a "final goods assembler" with a variety of open source or commercial components. Deciding what an atomic unit is in this realm is a very complex question.

Digital identity may be composed of a variety of accounts. For example, the combination of a GitHub account linked to a Twitter account and Google account might be a "digital identity" (Honan).

[9] Taken loosely. If they have a compromised IoT device, it may not have a user credential mapped to the person, but it could still be considered as "part" of the person in cyberspace.

The components, rigor, and risk associated with an identity can vary. A corporate identity may consist of a username, password, device profile, hard token, originating IP address, and time of day. In other contexts, the identity may be the device ID and default password on a consumer IoT device. In the latter case, the identity is of the device and not the user. The user may not be known to the device nor the system with which the device is interacting. New accounts can happen at the device level (unix userids), at a browser profile level, or within a session ("remember me on this device."[10])

Attackers are listed as compound elements because they're often working in a group or have a supply chain/market for their "product". Sometimes, these "threat actors" are denigrated as "script kiddies", and other times idolized as "APTs" or "ninjas".

Some of these units are inherently uncontroversial: people are people, accounts are accounts (with some possible nuance at the edges), and Google defines what a Google account is. Other units are less crisply defined or differentiated: is this malware sample the same as that one? Is this attack kit comprised of the same exploits as that?

Properties of these units are important. We usually believe that devices that are no longer getting updates are more vulnerable, and devices lose access to updates both because the manufacturer no longer creates those updates, and because some devices are re-configured to turn off updates. That reconfiguration is done both by owners, annoyed by updates or worried about them, or by attackers. So, measuring device age can be valuable. Device age probably correlates with wealth, and that may be a confounding variable. Wealthier people can both buy devices and have better access to support and help. One participant pointed out that software updates may carry problems, referring to the xz issue spreading as systems automatically updated. While this is typically called a "supply chain" issue, some participants assert that supply chains include relationships and are buttressed by contracts, SLAs, or warranties. Therefore, they take issue with the metaphor that open source software is part of a supply chain.

Some of the units may be politically sensitive; for example, capacity of factories or number of units sold.[11]

Research Question: Is the incidence of implanted problems growing? Knowing that requires us to track problems (the authors are not aware of a database). It also requires a consistent methodology for measuring the population of software.

[10] Which isn't even an accurate request — typically the prompt means "remember me in this browser."

[11] These are not universally sensitive; for example, Boeing and Airbus shout aircraft production numbers from the rooftops.

Research Question: What is the relationship between device age, wealth, and the observed security of devices?

Research Question: What data about units is gathered already (e.g. in import/export statistics)?

Properties about people may also matter. Similarly, there are assertions that either digital natives, elders with less experience/comfort with technology, or overworked parents are the most vulnerable or that young people pay the most attention or that wise elders are the most skeptical. So tracking age and other demographics of people or account holders may be useful. Each is accompanied by privacy concerns.

RELATIONSHIPS BETWEEN UNITS

We consider units assembled into populations; the public in public health. Public health often considers populations at scales like "the people of the United States," to measure the leading causes of death or life expectancy in a country. It will also frequently consider a subpopulation, especially one vulnerable to a disease. For example, uterine cancer rates are usually calculated for a population of women.

We can consider both subpopulations of humans and of computers in Cyber Public Health. Returning to the sexting blackmail, what is the right human population to count? Unmarried people? Should we count married people who might be looking for something either on the side or to spice up a committed relationship? On the computer side, if malware tends to infect a given platform, it seems useful to know those populations.

This brings us to the question of ontology. Some participants felt that differentiating factors in ontology may come down to how well defined they are and how well the relationships are mapped relative to one another. Others felt that more than one ontology may helpfully support or reflect the complexity of the underlying systems.

There are aspects of cybersecurity that are hard to quantify:

- Data
- Services

When talking about breaches of personal information, data is sometimes quantified by records, but the content of a record can vary from a name or a postal address to a credit card number or social security number or medical record. These gradations cover intrusiveness, difficulty of replacing, and data size. Data size is not a direct indicator of impact. A few bytes changed in a genomic dataset might obscure a crucial gene and its impacts. Data on psychographics and preferences are being gathered by more and more entities, and many of us would prefer not to have either our actual or

inferred preferences publicized or traded.[12] There are doubtless other elements, and a great deal of data is not about people, but its security is important to its owner. As storage becomes ever cheaper, measuring data by size seems less and less useful.

A participant suggested that data and service are broader than security, but fundamental units of computing, which raises the interesting specter of scope.

Research Question: How else might data be quantified? Are there types, valuations, or other measures?

"Critical infrastructure sector" is an important compound unit in American policy. It draws out the importance of system-oriented or socio-technical understandings of compound units. It may be that a public health framing can help with those complex conversations.

ACTORS, VICTIMS, AND THE PUBLIC

There was discussion of who to focus on. In security, it's common to focus on actors who *cause* harm and/or actors who *experience* harm. These are magnified in the criminal and conflict frames. In public health, there's a tendency to focus on sufferers (and those at risk/susceptible). There are also people who focus on pathogens or poisons; are they included in the public health world or the medical world?

In our imagination, a preponderance of health dangers is from non-sentient agents, primarily pathogens. An important result of the workshop is that public health does consider risks that result from human behavior: consider sexually transmitted diseases, smoking (enabled by companies who profit), or the use of motorcycle helmets (often regulated). It's worth noting that each of those human behaviors is a health concern because of their impact on the health of people. Cyber bullying is a concern for traditional public health because of its impact on mental health and suicide. Similarly, social media use drives anorexia in teenage girls (Dane).

How does human behavior relate to attackers in cybersecurity? There are ways in which we, through choices (freely made or forced), increase our vulnerability. In cybersecurity practice, we may choose to not do threat modeling before we design and build software, we may not do red teaming, we may not implement secure defaults, we may not enable auto-updating and patching, we may choose to ignore vulnerability reports, we delay patching, or we implement security and IT products poorly. In some cases, these are deliberate choices, and others are borne of ignorance. The result, nevertheless, is one of increased vulnerability to cyber exploitation. These behaviors, and the choices we make with regard to them, impact Cyber Public Health, and with increasing interdependence among technologies, these singular decisions can have significant collective impact.

[12] Or even gathered - a privacy issue- but once gathered, keeping that data confidential becomes a matter of security.

Research Question: What is the role of human choices in public health and how are they characterized? How can this inform Cyber Public Health and enable research to rapidly draw on a framing? What differences exist, for example, because of the technological basis of threats and vulnerabilities?

In medicine and public health, we have standardized lab/test result values (e.g. A1C has acceptable values and anything outside of that reference range may indicate disease). What are the cyber equivalents? Suggested possibilities include:

- Open Cybersecurity Schema Framework, mentioned above, helps with log management.
- Common Vulnerabilities and Exposures (CVE) system
- We might have an equivalent agent (APTs, UNCs, others) but companies have different names for it so we may not know. It's as if there were different classification schemes for illnesses and medication in the public health domain.

Assessing harms can be challenging, even where it may appear simple on the surface. An example is malware.

There are gradations of software that harms people. At one end is ransomware, clearly harmful to people. Microsoft defines a category of "Potentially Unwanted Apps (PUA)" and says they "are a category of software that can cause your machine to run slowly, display unexpected ads, or at worst, install other software that might be unexpected or unwanted. PUA isn't considered a virus, malware, or other type of threat..." (Microsoft, "Detect and block"). These are sometimes offered by actual companies with actual lawyers who think that their software being labeled malware is defamatory. TikTok has been accused of misbehavior with data they collect, and Facebook is being either sued or prosecuted for violations of the Wiretap act (Zeff).[13] These gradations and conflicts are relevant to Cyber Public Health because if we're attempting to measure harm, a disagreement between collectors about how to categorize TikTok could lead to dramatically different measurements.

We could measure harms by looking at type (theft of information, encryption of information, slowing a machine by mining bitcoin) and duration of the impact. There's an analogy of a COVID vaccine: a percentage of people were impacted temporarily by symptoms that caused them to lose a day of work and experience distress, but were provided longer term immunological benefit. A software agent, running on an endpoint to measure the health of that device, imposes an operational tax on the device in exchange for larger benefit to both the device and the population of devices under management.

[13] The cited case is confusing; Gizmodo links to https://www.courtlistener.com/docket/18714274/klein-v-meta-platforms-inc/ which is a civil action; the answer seems irrelevant for this report. Whoever is on the other side of the case from Facebook certainly believes Facebook's actions crossed a line.

We can consider analogies to tobacco companies, and perhaps there's a good paper there.

There is also an issue of costs. While human disease does have financial costs, including the cost of treatment or lost wages, and death clearly carries a loss of income, we also recognize that much of the cost of disease is emotional. In contrast, in Cyber Public Health, the costs are primarily financial. Many costs are the result of crimes, such as ransomware. There was discussion of both direct and reputational costs but, fortunately, not a single one of our wise participants brought up the discredited claim of stock price impact.

Insurance companies were repeatedly invoked as a possible data source, and it may be worth considering anonymized reporting requirements as part of a national insurance backstop. Other participants wryly pointed out that insurance companies seem to have market barriers to collecting data as part of onboarding new customers, and culpability barriers to data collection as part of claim resolution. It's possible that ISACs or ISAOs could be a third-party safe zone[14].

At a very different end of the spectrum, some participants discussed the difference between operational issues and strategic issues.[15]

DEFENSIVE ACTIVITY

One participant brought up five different levels of analysis that could be considered[16]:

- Asset
- Dyadic actor
- Systems: the entire Internet as a whole
- Strategic consideration
- Operational

This ties to questions of "is defense getting better faster than offense?" and "what works at scale?" and invokes Dan Geer's rubric of "no silent failures" (Janofsky).

Research Question: Are there public health measurements of defense (perhaps stockpile of PPE, readiness to produce vaccines) that inspire new research in cybersecurity?

RESILIENCY

[14] These are Information Sharing and Analysis Centers and Organizations; each is established with specific goals in U.S. law.

[15] An analog to individual disease vs. pandemics seems different. Perhaps there's interesting insight in considering either lifespan or quality of life measures and how they relate to cybersecurity issues at the strategic level.

[16] See (Healey).

Related to defense is investment in resiliency. We have a health system that is essentially responsive and recovery-oriented (e.g. ERs, acute care hospitals, insurance payment based on procedures done rather than procedures avoided), and that has also tended to be applied in our cyber health behaviors. We spend more on response and recovery on, say, ransomware, than on prevention activities. Often, post-breach, there is an increase of company spend on security.

But we know that events will probably happen, and so we can do things like buy Band-Aids and Tylenol in advance, build and staff ERs, and invest in 911 services. Likewise, in security, we can architect for failover, have backups and tested restoration, IR retainers, tabletop exercises, and more. Consumers can back up to iCloud, and some have second devices. We can measure the level of resiliency, perhaps by surveying or partnering with ASPR and FIRST and surveying industry and consumers.

- How can we measure cyber resiliency and recovery investment?
    - Units of measurement: lives impaired, financial costs, time loss, last known good state/configuration.
    - Going deeper: How much time do you need, what skills do you need, what kind of personnel do you need, how will the supply chain affect you?
    - What are the hazards against which we should be resilient? In the physical world, we seem to have global flu pandemics about every 100 years, and other epidemics (AIDS) are irregular. Some experts have suggested they will be more frequent. Beyond public health, we have hurricanes, fire, and flood.

- And then a different perspective is looking at this across a timeline:
    - Preparedness → response (during event) → recovery (post event)
    - We could measure time in phases. Some people are publishing measures of "dwell time" or "time to detect."
    - Those measures are focused on enterprises. Is anyone measuring cost to individuals, for example, time spent or what fraction of Geek Squad calls are security remediation?

Research Question: Some harms in public health are seasonal/cyclical (e.g. the flu) - is there a Cyber Public Health equivalent? For example, does spam rise during tax season, or does the level stay the same with a change in focus? Is the answer uniform across countries?

Research Question: What data do we collect in time series or could we collect in time series?

Research Question: How resilient are companies in the face of a cybersecurity event?

- How many organizations have defined a recovery time objective (RTO) of when they can reestablish digital operations, even in a diminished capacity?

- How many companies with these defined targets actually achieve them when tested or attacked?
- How many have also defined recovery point objectives (RPO), and know the amount of tolerable data loss in the event of a security incident? How often is that target achieved in testing or in an event?
- How many organizations have established plans for business continuity (BCP) or disaster recovery (DRP) - the equivalent of getting preventive vaccinations?
- How many organizations know their current level of vulnerability? How many are within a target level of vulnerability and patching?
- How many individuals, companies, or countries know how long they could and what they could operate without access to the internet?

Research Question: What are the hazards that we care about in Cyber Public Health?

- Issues such as heartbleed, log4shell, and shellshock (internet-facing, easily exploited, pre-auth vulnerabilities) seem to happen every few years
- Could we create an equivalent of 10-, 50-, and 100-year floods to help with resiliency?
- How do we think about hazards that require replacing hardware?[17]

2.3 DATA SOURCES

Having considered the harms that can come to various elements of the system, we can assess where we might get data on the units or harms to them that have been outlined above.

Perhaps most provocatively, an argument was put forth that "the data is out there, we would just need someone to pay for it." It's not clear that this is true, but it is worthy of investigation.

Research Question: What is the availability of data to answer the research questions in this report? Who can provide that data?

- Would the available data be "universally" accepted? Would it be challenged, questioned, or seen as unusable?

Research Question: Are there research questions where we know data is not available?

- e.g. non-customer impacting breaches of privately held companies.

Research Question: Are there research questions where data would be controversial to use?

- e.g. victim lists published by ransomware operators; using data from criminals about their victims raises ethical concerns.

[17] For example, an attack on Saudi Aramco destroyed the operating system of 30,000 machines and they needed to be manually re-installed. What would have been the impact if Aramco needed to buy 30,000 new machines?

Research Question: Public health authorities around the world do food inspections from "farm to table," at least for restaurant tables. What can we learn from that system of systems?

Research Question: Can someone make money by collecting and selling data?

Research Question: How can we quantify the value of a centralized, data gathering "bureau"?

We identified the importance of transparency about collection methods as well as the challenges introduced by a lack of alignment of collection methods across data sources. Additionally, it's important that the definitions are clear; the discussion of malware elsewhere shows that if one data gatherer calls TikTok malware, while another doesn't, their results will differ dramatically.

There was extensive discussion around entities that could collect data. We enumerated some potential data providers and types of data:

- Law enforcement
- Third parties (systems/services being provided)
- Commercial/non-commercial (data reviewed and scanned for vulnerabilities)
- Govt agencies (i.e. regulators)
- Other groups (BitSight, Security Scorecard)

Reporting requirements were a topic of conversation. Doctors are obligated to report certain diseases and, in parts of the U.S., they're also required to report gunshot wounds. Not doing so can result in a loss of license. Similarly, there are reporting requirements for labs, for device makers, and possibly others. In aviation, pilots, mechanics, and others are encouraged to report near misses. Part of the incentive is a reduction in penalties including loss of license. We do not have such a licensing regime in cybersecurity.[18]

Industry regularly opposes reporting requirements, a topic which we did not delve into in the workshop. However, there are transparency and reporting success stories, which are not routinely juxtaposed or analyzed with such resistance, including certificate transparency and VirusTotal (both of which were brought up as success stories). A great deal of information is shared with law enforcement; perhaps there's a way to get systematic output.

Research Question: What can the success of VirusTotal or Certificate Transparency teach us? What factors underlie those? Did they overcome resistance, avoid it, or present a different balance?

[18] Section 2.5 of The ACM Code of Ethics states "Computing professionals are in a position of trust, and therefore have a special responsibility to provide objective, credible evaluations and testimony to employers, employees, clients, users, and *the public*... Any issues that might result in major risk must be reported to appropriate parties."

Research Question: How much does law enforcement report as a fraction of what they learn about (for various law enforcement agencies)? How systematic are the crime reports which are produced? Do they collect information that might be useful? The success of "accident report forms" may be relevant (99% Invisible).

In the reporting requirements discussion, it was suggested that we could explore compelled disclosure (enterprise-level) of incident(s) after 12 months, or reporting-by-default on every device to enforce mandatory device posture reporting to a central agency. This would allow individuals to report their loss of their Digital ADLs.

It was suggested that major cloud providers have insight into issues with their customers. For example, malware command-and-control connections, or spikes in compute due to cryptocurrency mining or storage access because of ransomware. Could a liability shield for reporting, possibly anonymized, be helpful?

3. DIRECTIONS

Discussion about "what comes next?" was enthusiastic, and we've identified three main themes: research, advocacy, and impact. Research is about academic research. How do we build a community, develop venues in which Cyber Public Health work can be refereed, shepherded, and shared? How do we get funding for interdisciplinary work? Advocacy is about getting these ideas into the public conversation about cybersecurity. Lastly, impact is about finding and celebrating success stories that inform each of those. We'll address them here in reverse order.

3.1 IMPACT AND SUCCESS STORIES

The history of public health is full of stories of individuals who chose to make a difference. From the relatively well-known story of John Snow and the Broad Street pump to more obscure stories like Hunziker's introduction of iodine in Switzerland (Goodman). These sorts of stories likely informed funding for things like the British Doctors Study and the Framingham Heart Study, each of which contributed to our knowledge of the dangers of tobacco. We could go on, but rather we'll turn our attention to the importance of *demonstrable success* in the journey, and comment that nothing succeeds like success. Metaphors certainly don't, nor do exhortations. We have plenty of medical and public health metaphors -- we gain when we show that the metaphors can teach us things, drive effective interventions or illustrate that problems are better or worse than we thought.

The stories are backed by data and evidence of success -- John Snow and Reverend Whitehead did create a map, they did get the handle removed from the pump, and they did stop that cholera outbreak (Johnson). It then fostered a larger set of actions leading to modern understanding of disease and epidemiology.

We are looking to find similar stories in Cyber Public Health, where there were actions taken that changed the behavior of people or systems to be healthier. One participant mentioned Microsoft and the work done with Windows XP SP2 and Windows Update Service, which introduced the ability to assess and distribute needed patches to systems. Prior to this, it required user action to go to a site and download a specific patch for a system and apply it, or install from a CD-ROM. Google Chrome can push updates to the browser continuously without user action and implement via restart with full tab restoration. These are, perhaps, small examples of Cyber Public Health interventions that are addressing a class of dangers. Other examples include Zero Trust that moved trust away from network-based authentication.

Vital statistics are a possibility for early successes. Measuring births, deaths, or current population will involve working through challenges like defining that population, and the ontological/taxonomic issues of defining equivalents to births or deaths.[19]

Use cases seem like another possible early success. Defining how data could be used or what specific problems could be solved for whom may make this a reality. One participant noted, "Policymakers need a broad scope. They shouldn't approach it in the same way as a vendor pushing out a security update."

Research Question: How many of the National Cyber Performance Goals could be addressed with Cyber Public Health data?

Research Question: What incentives are there for incident reporting? Could we get, for example, Blackhat or Usenix Security or CISA to present an award for transparent and helpful breach notification?

Research Question: What are the sources of political opposition to a field of Cyber Public Health and are there frames or values that we could deliver to get them to see that Cyber Public Health's costs are worth paying?

3.2 ADVOCACY

There are many communities that might benefit from and contribute to Cyber Public Health.

For example, the United States government (at the level of the National Cyber Director) looks to measure the impact of the National Cybersecurity Strategy. Many of the questions within the strategy could be answered better with vital statistics. Someone could expand that sentence into a briefing paper, and someone could bring that paper to the right folks.

[19] CyberGreen is currently pursuing some grants in this area.

Many of the ideas we are working with may be interesting to "traditional" public health, to the privacy community, to the 'trust and safety' community, to communities studying bullying or mental health online; some of their work is interesting to us. We should engage with communities such as the Privacy Law Scholars, the University of Washington's Institute for Health Metrics and Evaluation or the Georgetown School of Public Health.[20]

Tools like "easy onboarding" documents that introduce the ideas behind Cyber Public Health, community spaces for discussion, and more were discussed. This workshop report will be one such document.

Eventually, institutions will need to gather, analyze, and publish data; drawing attention to the need for institutions which parallel those functions of CDC and WHO is work that was identified at the workshop. The Solarium Commission-proposed Bureau of Cyber Statistics may be an early such institution; we should consider further advocating and pushing for its development through the framework of Cyber Public Health.. One group asserted that public-private partnerships for sourcing/reporting on cybersecurity incidents are *essential* and should be pursued.

Discussion touched on "professionalization". Both research and advocacy would be needed to create the structures of professionalization including exams, licensing, and professional bodies.

Research Question: What scope would be professionalized? Computing? Cyber Public Health? Cyber Epidemiology?

- Would there be licensing? Licensing is generally imposed either to protect the public or limit trade, and there are arguments that licensing is a form of restraint of trade. How would those play here?
- There are arguments for and against computing as a profession; what do they tell us about this?[21]

Research Question: How can incentives be used to facilitate improvement?

- Is there a possibility of "cyber offsets" like we have with climate/energy efficiency?
- Service providers could also do it on an organization's behalf in the same way that HR Block does your taxes for "free" but knows a certain percentage of the population will get a return and HR Block will keep a percentage of that return.
- ISPs/CSPs/other service providers (e.g. Google for GMail) should have an incentive to report out information. We already see something similar with credit card fraud reporting, Google's Transparency Reports, and Exposure Notifications for COVID.

[20] This is not an exhaustive list, just some examples.

[21] The argument is rather extensive and examples include Chien, Choi, Denning, National Research Council.

3.3 RESEARCH

Moving from a metaphor to a discipline will involve showing that the work has value, identifying questions which align with the research goals of various communities and funders, and creating venues for such work to be published and found. We established such venues with this workshop and with the CyberGreen/Ostrom recurring workshop series. We are also exploring ideas such as a Research Coordination Network grant and engagement with the Computing Research Association.

There was discussion of the value of a research agenda and a shared bibliography. (We have started a shared bibliography in Zotero.[22]) There was brainstorming around possible research projects that seem "within reach" (i.e. research that could be undertaken using approaches that don't require statistical institutions to be in place). Such approaches could include telephone/internet surveys, observational studies, and diary or journal research. Research questions of possible projects include:

- How much time does a person spend on security activity over a month (including patching, password resets, waiting on SMS/email MFA, completing Captchas, etc.) How does this differ across professions, countries?
- How frequently are Digital Activities of Daily Living interrupted by security toil that takes more than 15 minutes to resolve?
- How much e-waste is driven by unfixable security problems? For example, what's the e-waste impact of Windows 11's new hardware requirements for security?
- Can we create more standardized incident reporting forms? This seems relevant to concerns about growing costs of incident reporting regulation.
- What would a cybersecurity mortality and morbidity report cover? Can we create mockups? There is a sense that it would differ from today's industry reports, specifically including the Verizon DBIR. Can we make that more concrete?
- How can we get to predictive capability?
    - Could we create a risk-o-meter that provides an individual with a view into their behavior? For example, "Today, your behavior is riskier than usual and increases the chance that you'll be subject to malware/cyber risk."[23]
    - Can we predict incidence within a population? Can we measure incidence within, say, critical infrastructure sectors?[24]
- How do the social and technical aspects of security (i.e. socio-technical security) come together and interact?

With help from partners, we could also investigate:

[22] https://www.zotero.org/groups/5355707/cyberpublichealth

[23] Investigation into what data, specifically, such a meter would require and why would advance our specific understanding.

[24] This seems within reach in collaboration with the Sector Risk Management Agencies; and we could use Delphi methods to predict the year on year changes given regulatory changes.

- What's the distribution of recovery timelines, and how does that relate to criticality (e.g. work/school accounts vs. "personal" accounts vs. bank accounts vs. software systems that are work critical.)
- If VirusTotal took a public health perspective, what could we learn?
    - One idea was to take inspiration from a project where individuals submitted their Explanations of Benefits and it showed how your benefits could apply differently depending on the provider you saw; this led to price transparency.

Over time, maturation towards predictive capability to incentivize change in operations as a preventative measure could help to guard against future cybersecurity risk. This maturation will also enable us to test theories and demonstrate value, although we should not ignore the Cassandra Effect or the Heisenberg Effect.[25]

Our workshop focused primarily on cybersecurity and public health. It seems that a lot of today's AI measures are focused on instances of vulnerabilities or incidents, with limited measure of population impact. It seems reasonable to think that a public health approach could apply to AI, enumerating human-centered harms and creating population-centered measurements. For example, how many people believe that their job search has been hindered by AI? How many people are choosing to not use LLMs to help them get their jobs done because of concerns about hallucinations?

3.4 OTHER POSSIBLE DIRECTIONS

There was discussion around potential next steps and solutions which include professionalizing the industry, reducing liability related to reporting, incentivizing good practices, and implementing government regulations.

Some challenges around professionalization have been discussed above. One additional aspect to consider with respect to professionalization is education. What would be in a Cyber Public Health course or curricula?

One group listed "Get consensus on basic units of interest because absent this, there's a lack of clarity on what should be measured. This includes our understanding of the targets – Actors? Sufferers? Both?" This brings up an important question about the value and prioritization of consensus. Public health has very clearly defined terms like incidence, and that clear definition is a result of consensus. It may be that developing consensus enables faster research because data is gathered and expressed in ways that reflect that consensus, or that a consensus process becomes interminable and focuses on definitions that are not useful.

[25] The Cassandra Effect and the Heisenberg Effect, respectively, are the contradictory ideas of being able to predict the future but being ignored, or that our predictions will alter the thing we're measuring.

Research Question: What can we learn from other sciences, disciplines, or standards processes and their consensus processes?

Research Question: What research goals depend on consensus, and which ones can be pursued in its absence?

Research Question: What is the cost/danger/risk of delaying consensus processes?

4. CHALLENGES

We can anticipate many challenges and expect to be surprised by others.

4.1 PRIVACY

Some of the data we'll gather will be at a population level; other data will be at least personally identifiable (e.g. usernames) and possibly sensitive (e.g. how they got compromised). Being aware of and sensitive to these concerns will be important and will require data governance. At least one paper, Public Health as a Model for Cybersecurity Information Sharing (Sedenberg and Mulligan), has considered some of these issues. Cryptographic techniques including but not limited to Zero Knowledge techniques and Differential Privacy may allow us to manage some of these; others will require attention to contextual integrity and new norms. Other techniques including federated learning and confidential computing may be useful.

Privacy advocates may also benefit from our work: Issues impacting Digital Activities of Daily Living probably include both privacy concerns and incidents. Privacy concerns leads to reluctance to engage or disengaging from previous activities, and incidents lead to both harms and responses.

There may be data which businesses wish to keep secret, including population and incident data. Christopher Morten has written on Publicizing Corporate Secrets.

4.2 INTERDISCIPLINARITY

Academic disciplines are a powerful organizing concept. Interdisciplinary work brings challenges at every stage. Who funds it, what venues will publish work, and who can review as a peer? What courses will teach the work and excite future students? But bringing multiple disciplines together is often a powerful way to help new ideas flourish.

5. NEXT STEPS

There was general agreement that subsequent workshops around Cyber Public Health would be an effective way to follow up and a way to continue to push the research and advocacy agendas. In

addition, CyberGreen Institute has co-founded a Cyber Public Health working group and is managing a mailing list to keep interested parties abreast of its activities. This includes the aforementioned (~monthly) CyberGreen/Ostrom recurring workshop series.

As organizers, we'll aim to continue to host workshops and write reports like these on an annual basis, provided we get the requisite funding. We seek to foster and broaden our community, and we ask our current community members to spread the word.

6. PARTICIPANTS

The workshop was operated under the Chatham House Rule, which is formally that "participants are free to use the information received, but neither the identity nor the affiliation of the speaker(s), nor that of any other participant, may be revealed." We are respecting the first two parts of the rule while listing attendees by agreement. We offered an opt-out option for any participant who chose to not be identified. Participant affiliations are not listed to avoid any mis-attribution, and no statement in this report should be attributed to any participant.

Carol Bernard
Whitney Bowman-Zatzkin
Jean Camp
Yi Ting Chua
Jane Coffin
David Conrad
Josiah Dykstra
Charles Fracchia
Dan Geer
Jason Healey
Trey Herr
Matt Hladnick
Douglas Hough
Yurie Ito
Marina Kaganovich
Rob Knake
Taylor Lehmann
Genevieve Liveley
Art Manion
Luke McNamara
Thomas Millar
Grant Ongers
Rob Reeder
Bill Reid
Phil Reitinger
Danielle Ruderman
Blake Scott
Scott Shackleford
Adam Shostack
Jared Smith
Simon Sun
Nathan Taback
Arastoo Taslim
Stan Trepetin
Luis Urena
Kris Yun

7. BIBLIOGRAPHY

99% Invisible. "The Nut Behind the Wheel." 5 Dec. 2017, https://99percentinvisible.org/episode/nut-behind-wheel/.

"ACM Code of Ethics." ACM, 2018, DOI: 10.1145/327459, https://www.acm.org/binaries/content/assets/about/acm-code-of-ethics-booklet.pdf.

Charney, Scott. "Collective Defense: Applying the Public-Health Model to the Internet." IEEE Security & Privacy, vol. 10, no. 2, 2012, pp. 54–59, https://ieeexplore.ieee.org/document/6051417.

Chien, Andrew A. "Computing Is a Profession." Communications of the ACM, vol. 60, no. 10, 2017, p. 5, https://dl.acm.org/doi/fullHtml/10.1145/3137136.

Choi, Bryan H. "Software Professionals, Malpractice Law, and Codes of Ethics." Communications of the ACM, vol. 64, no. 5, 2021, pp. 22-24, https://dl.acm.org/doi/fullHtml/10.1145/3457193.

CyberGreen. "The Cyber Green Initiative: Concept Paper Improving Health through Measurement and Mitigation: Cyber Green Initiative Concept Paper." 2014, https://www.jpcert.or.jp/research/GreenConcept-20141117_en.pdf.

Cohen, Fredrick. "Computer Viruses: Theory and Experiments." 7th DOD/NBS Computers & Security Conference, Sept. 1984.

Dane, Alexandra. "The Social Media Diet: A Scoping Review to Investigate the Association Between Social Media, Body Image, and Eating Disorders Amongst Young People." PLOS Global Public Health, vol. 3, no. 3, 22 Mar. 2023, e0001091, https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10032524/.

Denning, Peter J. "The Computing Profession." Communications of the ACM, vol. 61, no. 3, 2018, https://cacm.acm.org/opinion/the-computing-profession/.

Edemekong, Peter F., et al. "Activities of Daily Living." 2019, https://www.ncbi.nlm.nih.gov/books/NBK470404/.

Goodman, Jonah. "A National Evil." London Review of Books, vol. 45, no. 23, 30 Nov. 2023, https://www.lrb.co.uk/the-paper/v45/n23/jonah-goodman/a-national-evil.

Hansen, Lene, and Helen Nissenbaum. "Digital Disaster, Cyber Security, and the Copenhagen School." International Studies Quarterly, vol. 53, no. 4, Dec. 2009, pp. 1155–1175, https://doi.org/10.1111/j.1468-2478.2009.00572.x.

Healey, Jason. "Understanding the Offense's Systemwide Advantage in Cyberspace." Lawfare, 22

Dec. 2021, https://www.lawfaremedia.org/article/understanding-offenses-systemwide-advantage-cyberspace.

Honan, Mat. "How Apple and Amazon Security Flaws Led to My Epic Hacking." Wired, 6 Aug. 2012, https://www.wired.com/2012/08/apple-amazon-mat-honan-hacking/.

Janofsky, Adam. "Cybersecurity Guru Dan Geer on Quantum Computing, the Future of Security, and Running a Farm." The Record, 12 Mar. 2021, https://therecord.media/cybersecurity-guru-dan-geer-on-quantum-computing-the-future-of-security-and-running-a-farm.

Johnson, Steven. *The Ghost Map: The Story of London's Most Terrifying Epidemic--and How It Changed Science, Cities, and the Modern World*. Penguin, 2006.

Knake, Robert, Adam Shostack, and Tarah Wheeler. "Learning from Cyber Incidents: Adapting Aviation Safety Models to Cybersecurity." Harvard Kennedy School Belfer Center, 12 Nov. 2021, https://www.belfercenter.org/publication/learning-cyber-incidents-adapting-aviation-safety-models-cybersecurity.

Microsoft. "Detect and Block Potentially Unwanted Applications." 28 Aug. 2023, https://learn.microsoft.com/en-us/microsoft-365/security/defender-endpoint/detect-block-potentially-unwanted-apps-microsoft-defender-antivirus?view=o365-worldwide. Accessed 28 Mar. 2024.

---. "Zeroing in on Malware Propagation Methods." Microsoft Security Intelligence Report, vol. 11, 2011, https://download.microsoft.com/download/0/3/3/0331766E-3FC4-44E5-B1CA-2BDEB58211B8/Microsoft_Security_Intelligence_Report_volume_11_Zeroing_in_on_Malware_Propagation_Methods_English.pdf.

MITRE. "Common Malware Enumeration List." MITRE, https://cme.mitre.org/data/list.html. Accessed May 2024.

Morten, Christopher J. "Publicizing Corporate Secrets." University of Pennsylvania Law Review, vol. 171, 2022, p. 1319, https://scholarship.law.upenn.edu/cgi/viewcontent.cgi?article=9816&context=penn_law_review.

National Research Council, et al. *Professionalizing the Nation's Cybersecurity Workforce?: Criteria for Decision-Making*. National Academies Press, 2013, www.nap.edu/openbook.php?record_id=18446.

Peng, Roger D., and Elizabeth Matsui. *The Art of Data Science*. Bookdown, 2017,

https://bookdown.org/rdpeng/artofdatascience/.

Sedenberg, Elaine, and Deirdre Mulligan. "Public Health as a Model for Cybersecurity Information Sharing." Berkeley Technology Law Journal, vol. 30, no. 3, 2016, pp. 1687–1740, http://dx.doi.org/10.15779/Z38PZ61.

Slupska, Julia, and Mariarosaria Taddeo. "Generative Metaphors in Cybersecurity Governance." *The 2019 Yearbook of the Digital Ethics Lab*, edited by Christopher Burr and Stefania Milano, Springer, 2020, https://doi.org/10.1007/978-3-030-29145-7_2.

Shostack, Adam. "Cyber Vital Statistics." CyberGreen Technical Report, TR22-02, 2022, https://cybergreen.net/technical-report-22-02/.

Solove, Daniel J. *Understanding Privacy*. Harvard University Press, 2010.

Wash, Rick. "Folk Models of Home Computer Security." Proceedings of the Sixth Symposium on Usable Privacy and Security, 2010, pp. 1-16.

Wolff, Josephine. "Cybersecurity as Metaphor: Policy and Defense Implications of Computer Security Metaphors." 2014 TPRC Conference Paper, Mar. 2014.

Zeff, Maxwell. "Project Ghostbusters: Facebook Accused of Using Your Phone to Wiretap Snapchat." Gizmodo, 26 Mar. 2024, https://gizmodo.com/project-ghostbusters-facebook-meta-wiretap-snapchat-1851366093.