Columbia World Projects
Cybersecurity Forum Report

December 14, 2018

Foreword

Dear Reader,

On behalf of Columbia World Projects (CWP), we are pleased to present the following report on the results of our Forum on Cybersecurity, one of a series of meetings we are holding to identify possible solutions to fundamental challenges facing humanity.

Cybersecurity – the practice of protecting systems, networks, programs, and data from digital attacks and unauthorized access – has grown increasingly challenging as a consequence of increased connectivity through the Internet. It is hard to think of a tool in our contemporary society that brings with it such potential for good and for ill as the Internet. On the one hand, it can be a catalyst for economic growth and for social change, and can expose us to a wide range of information and perspectives. On the other hand, it can be used to wreak havoc on financial and commercial networks, to repress dissent, and to undermine democratic processes. The Internet is essential to our personal lives and to our jobs, from communicating with our colleagues, to receiving payments for the goods and services we provide. Yet the very qualities that make the Internet so valuable to society, including the immediate interconnectedness it facilitates in a relatively decentralized and open way – make the task of securing the services, devices, data, and infrastructure that enable these key functions extraordinarily challenging.

It is commonly accepted that efforts to identify, understand, and address the risks – which continue to evolve and multiply with our growing reliance on the Internet – have been insufficient, and too often reactive rather than proactive. The same can be said for our response to the increasingly complex threats to standalone systems, networks, and programs. Attempts to manage cybersecurity have been constrained by existing and sometimes outdated organizational and institutional frameworks, outpaced by technological advancements employed by malicious actors, and subject to the significant challenge of collaborating across the public and private sectors. Furthermore, the rapid development and evolution of new technologies, such as Artificial Intelligence and machine learning, continue to create new ways to both carry out cyber attacks and defend against them. Yet we are not doing enough to mitigate the former or to take advantage of the latter.

For all of these reasons, we decided to focus a CWP Forum on cybersecurity. On September 25, 2018, more than 35 experts from inside and outside of Columbia University – representing a range of substantive and institutional perspectives – came together at Columbia in the City of New York. The aim of the meeting was not only to deepen our understanding of complex cybersecurity challenges, but also to identify a set of promising projects in which Columbia faculty and researchers could partner with practitioners to implement potential solutions. These project proposals, developed in advance of the Forum through a collaboration between experts and CWP staff, were aimed at finding pragmatic ways to reduce the profound vulnerabilities and threats inherent in the Internet – and more broadly in the digital sphere –

without undermining its capacity for good. The attached report represents the work that took place at that Forum.

While a list of the experts who participated in the Forum and helped draft the report is included at the end, the ideas and views it contains are not attributed to individual participants or organizations, as was agreed in advance of the meeting. Yet the report does try, where possible, to specify the relative support for an idea or point of view, ranging from an individual observation to a consensus view.

Even as identifying project ideas to be developed by CWP and its partners is the primary objective of this and future Fora, it is also our objective that these gatherings will deepen the understanding of complex global challenges, inspire even the most advanced experts to see vexing problems in new ways, and encourage partnerships that might lead to breakthroughs that improve lives. Thus, in sharing the ideas and insights of experts who have generously given their time and intellectual capital to our effort, we hope others will benefit from the conclusions they reached and share their own thoughts on these matters with CWP, as we continue to seek ways to effectively tackle these challenges.

In a similar vein, there are some ideas that were developed for the Forum that, while not the right fit for further development by CWP, are ones we believe are worth pursuing. In those instances, we are working to foster partnerships and open pathways, both inside and outside of Columbia, that will allow these ideas to continue to develop. We know that complex challenges like these cannot be solved by any organization or institution alone, and that it will take many efforts to make meaningful progress.

Nicholas Lemann
Director, Columbia World Projects

Avril Haines
Deputy Director, Columbia World Projects

Nik Steinberg
Forum Director, Columbia World Projects

Table of Contents

I. Introducing the Challenge

The Internet is integral to our daily existence and often has a positive impact on our work, play, and personal lives. It allows us to connect instantly with friends, family, and colleagues, and helps us build and participate in new communities. It gives us greater access to knowledge, markets, entertainment, and ideas, and countless other resources than at any time in human history. Digital tools promote innovations through global networks, and allow people to share medical and safety resources and support swiftly in times of crisis. The Internet has provided access to education where it was not previously available, and has enabled improvements in the quality of education where it was available. It is where more and more people go to learn what is happening in the world on any given day, and where we engage in public debate. It has been integrated into every aspect of our national defense. We have even connected the cars we drive, the trains we ride, the planes in which we fly, and vital infrastructure serving our cities and towns, farms and factories to the Internet, in an effort to make them function more effectively. Yet while the value the Internet brings is clear, so are the myriad vulnerabilities that it creates.

Indeed, it is precisely because of the indispensable role the Internet plays in our lives that disruptions to the systems, and networks that undergird them, can rapidly bring so much of what we do to a standstill, undermine our privacy and civil liberties, and even threaten our prosperity and national security. And of course, it is not just systems that are connected to the Internet that are vulnerable to digital attack – so too are standalone systems, networks, and programs. While it is difficult to measure the precise level of exposure in this realm, there is a clear consensus that cybersecurity is one of the most significant and complex challenges facing the world today. To give just a few examples, the U.S. Director of National Intelligence said cybersecurity is one of his "greatest concerns and top priorities."[1] The U.S. Homeland Security Secretary's has assessed that cyberweapons and sophisticated hacking pose a greater threat to the United States than the risk of physical attacks.[2] Freedom House concluded that digital disinformation tactics have contributed to a global decline in Internet freedom every year for the last seven years, and played an important role in elections in at least 18 countries from 2016 to 2017 alone.[3] And estimates regarding the global economic impact of cyber attacks range from $400 billion to more than $2 trillion each year.[4] So too is there is a growing recognition

[1] Dan Coats, "Remarks as prepared for delivery by The Honorable Dan Coats Director of National Intelligence" (opening statement, Senate Select Committee on Intelligence, Annual Threat Assessment, February 13, 2018), https://www.dni.gov/index.php/newsroom/congressional-testimonies/item/1846-dni-coats-opening-statement-on-the-worldwide-threat-assessment.

[2] Kirstjen M. Nielsen, "Rethinking Homeland Security in an Age of Disruption" (speech, Washington, D.C., September 5, 2018), https://www.dhs.gov/news/2018/09/05/secretary-nielsen-remarks-rethinking-homeland-security-age-disruption.

[3] Freedom House, *Freedom on the Net 2017*, November 14, 2017, https://freedomhouse.org/report/freedom-net/freedom-net-2017.

[4] The true global economic impact of cyber attacks is unknown given, among other things, the lack of verifiable data and a common vocabulary. Yet even if one assumes that the most conservative estimate is accurate, the figure remains astounding. The figures cited above represent a range of estimates from 2015. See: Stephen Gandel,

that technology solutions alone cannot address the many vulnerabilities and possible vectors of attack, but rather that behavioral, normative, regulatory, social, and other interventions will also be critical to building effective solutions. Yet, despite these high-level warnings and the fact that a fair amount of attention and resources have been devoted in the last several years to cybersecurity, the increasing sophistication of cyber threats continues to outpace progress, as does the number of attacks, particularly in the United States and in Europe.

It was with this understanding – that current efforts to address cybersecurity are insufficient – that participants in the Columbia World Projects (CWP) Forum on Cybersecurity began their opening plenary discussion. Approximately 35 experts with a range of different substantive and institutional perspectives shared their views on the nature of the threat, key vulnerabilities, and the particularly intractable challenges associated with addressing them. This discussion provided critical context for the concrete project proposals taken up later in the working groups (Section II), and helped inform the selection of projects meriting further development by CWP (Section III).

The Nature of the Threat

The threat in cyberspace is often referred to as an asymmetric threat – one in which a major military power like the United States is vulnerable to attack by less powerful actors. In unpacking why this is the case, six aspects of the threat were highlighted by participants during the Forum's opening plenary session: (i) the low barrier to access by malicious actors – meaning that it is easy to access digitally, as opposed to physically, high-value assets through the Internet; (ii) the low cost of conducting an attack, particularly given that the tools for doing so are widely available, relatively inexpensive, and do not require significant expertise to use; (iii) the potential for significant damage and high-value return on an attack, given the degree to which we rely on the Internet and digitally-connected devices to conduct personal and professional activities, as well as to deliver public services that are critical to our survival; (iv) the degree to which it is possible to conduct an attack with impunity or even anonymity; (v) the asymmetric advantage that attackers have because they only need to find a single weak link in a network, while defenders need to secure entire networks; and (vi) the degree to which the interconnected and interdependent nature of digital networks means the consequences of an attack can spread almost instantaneously, with consequences for individuals and entities far beyond an initial target.

Participants also took note of the evolving nature of the cyber threat. One participant noted that attacks by foreign governments on the United States over the years have evolved from

---

"Lloyd's CEO: Cyber Attacks Cost Companies $400 Billion Every Year," *Fortune*, January 23, 2015, http://fortune.com/2015/01/23/cyber-attack-insurance-lloyds/; Steve Morgan, "Hackerpocalypse: A Cybercrime Revelation," *Cybersecurity Ventures*, August 26, 2016, https://cybersecurityventures.com/hackerpocalypse-original-cybercrime-report-2016/.

intrusions generally aimed at gathering intelligence, to intrusions aimed at influencing actors, with greater impact over time as the Internet has become more ubiquitous. Participants highlighted increasingly hostile activities on the Internet, principally by state actors, to promote disinformation; erode trust in facts; exacerbate existing divisions within societies; undermine the credibility of public institutions and processes; retaliate for perceived grievances (as in the 2014 Sony Pictures attack); steal intellectual property, including trade secrets or other confidential business information; collect private information for possible misuse; and conduct activities that make clear that an attack on a high value asset, such as critical infrastructure, could be deployed with little warning. Other participants called attention to the increasing number of intrusions by both state and non-state actors for financial gain, and of illegitimate invasions of privacy through intentional and unintentional disclosures of personal information.

Multiple participants pointed out that the fear of a massive cyber attack – what one participant called a kind of "cyber 9/11" – overshadows what in their view is the greater risk, which is the incremental harm inflicted by a multitude of smaller attacks that occur every day. By focusing disproportionately on preventing a single catastrophic attack, participants noted, we are failing to pay sufficient attention to stopping a "death by 1,000 hacks." Participants agreed, however, whether in the context of a catastrophic attack or numerous smaller attacks, there are key vulnerabilities to be focused on and protected. These include our **critical infrastructure** – systems and assets whose destruction or incapacity would have a debilitating impact on the national security, economic prosperity, health, or safety of our societies; and our **privacy**, particularly with respect to our personal information, which has become increasingly challenging to protect from both state and non-state actors. A third area of vulnerability that was discussed related to **information challenges** on the Internet, which includes not only efforts to manipulate or fabricate information, but also the way such information is shared and consumed online, especially on social media.

Finally, there was considerable debate among the participants on whether to include this last category within the scope of the Forum's discussion. While everyone agreed it was an important issue, some were of the view that it should not be included, as it was not strictly a cybersecurity issue. Others viewed the information challenges associated with the Internet as an existential threat to liberal democracies that could be impacted – at least in part – through digital or computational measures, and thus argued that the issue should be included in the Forum's discussion. As one expert noted, even determining what constitutes an information challenge on social media is challenging. While technology companies define information threats primarily in terms of external actors (such as foreign governments or violent extremist groups attempting to spread disinformation or propaganda), from a user's perspective, the companies themselves may present an information threat (for example, by privileging and amplifying speech that is false, incendiary, or corrosive to the public discourse). Another expert pointed out that approaches taken by social media platforms may make it easier (or harder) for adversaries to carry out information operations, and highlighted the importance of such

platforms for detecting and attributing anomalous behavior. Ultimately, we decided to include information challenges within the scope of the Forum.

Key Challenges

Participants discussed a range of intractable challenges that have made it especially difficult to address malicious activities in cyberspace. By far the most significant concerns related to organizational, behavioral, and institutional challenges associated with the various actors who have access to, and consequently must be part of securing, the Internet. While all agreed that technical developments were important to understanding the problem and to devising solutions, there was also consensus that the main challenges are not purely technological ones. In particular, participants underscored: (i) the importance of better collaboration, communication, and coordination among and between public and private sector actors; (ii) the need for behavioral changes and incentives that promote the use of basic cybersecurity practices by individuals, private companies, and public entities; and (iii) the importance of developing international norms and a strategy for deterrence, in order to promote stability and reduce the risk of inter-state conflicts sparked in cyberspace.

Collaboration, Communication, and Coordination Among and Between Sectors: Participants agreed on the importance of the private sector to any cybersecurity strategy. Roughly 85 percent of the United States' critical infrastructure is owned or operated by the private sector,[5] and according to a recent U.S. Department of Energy report, an estimated 90 percent of the United States' energy infrastructure is in the hands of the private sector.[6] Much of this infrastructure employs digital technologies and is connected to the Internet (e.g., the use of smart meters that promote more efficient, reliable, and cost-effective use of energy resources). Moreover, the private sector is essential to cybersecurity for reasons beyond its ownership of and control over critical infrastructure. Participants noted that supply chains that pass through the private sector are another important risk to be mitigated. Modern infrastructure and weapon systems, for example, are heavily dependent on microelectronics, which have a complex supply chain involving multiple industries, manufacturers, and distributors in a wide variety of countries. Recent work has underscored the need for new strategies to address this challenge.[7]

Participants also underscored that the government's ability to provide for the security of its citizens and lay the groundwork for economic prosperity is increasingly dependent on the

---

[5] U.S. Government Accountability Office, *The Department of Homeland Security's Critical Infrastructure Protection Cost-Benefit Report*, June 26, 2009, https://www.gao.gov/products/GAO-09-654R.

[6] U.S. Department of Energy, *Multiyear Plan for Energy Sector Cybersecurity*, March 2018, https://www.energy.gov/sites/prod/files/2018/05/f51/DOE%20Multiyear%20Plan%20for%20Energy%20Sector%20Cybersecurity%20_0.pdf.

[7] U.S. Department of Defense, *Report of the Defense Science Board Task Force on Cyber Supply Chain*, February 2017, https://www.acq.osd.mil/dsb/reports/2010s/1028953.pdf.

private sector to maintain the speed, resilience, and reliability of the Internet and related technology, as well as for critical information on these systems. As one participant noted, while in terrorism-related security challenges the government tends to possess far more intelligence than the private sector, in the realm of cybersecurity, the reverse is true: private sector entities often have more, and more up to date, information about specific threats and attacks, as well as about the different ways that systems are being compromised.[8]

In sum, there was unanimous agreement that neither the public nor the private sector could address cyber threats alone. Consequently, public-private sector collaboration, communication, and coordination on the cybersecurity standards to be applied in the construction and maintenance of relevant infrastructure – as well as methods and mechanisms for detecting, identifying, protecting, responding, and recovering from attacks – are critical to achieving cybersecurity. While particularly the case in the United States, this also holds for many other countries around the world.

In discussing the importance of integrated efforts across the public and private sectors, several participants highlighted the need for a clear allocation of responsibilities among major stakeholders within both sectors, pointing out that the current lack of agreement over roles and responsibilities has been a significant challenge in establishing better collaboration, communication, and coordination. Some participants indicated that the government and major Internet and platform providers should each be taking on greater responsibility for cybersecurity, and that a legal and regulatory structure that assigns such responsibility is needed. Others were of the view that a voluntary structure would be preferable. Nevertheless, all agreed that the current approach places too much trust in, and too great a burden on, individual users to make wise decisions and adopt best practices, when we know that individuals (and organizations, for that matter) on average do not prioritize basic cybersecurity practices. Experience has demonstrated time and again that this expectation is not only unrealistic, but dangerous, given that any weak link – all the way down to the level of an individual user – can be exploited to compromise an entire system. Participants said decision-making and responsibility for cyber vigilance should move away from the user end of the spectrum, and towards the public and private entities that possess a greater concentration of knowledge and resources. But in doing so, the lines of responsibility need to be clear, including the roles and responsibilities of different government agencies and departments focused on various aspects of cybersecurity, and they need to be communicated in such a way that they are understood by key stakeholders.

Participants also referred to the lack of real-time sharing of relevant, timely, and actionable threat information between and among critical infrastructure partners and relevant government entities, which could improve preventative and proactive response measures to

[8] Although, as one expert noted, the government tends to have more knowledge than the private sector regarding the malign efforts of state actors, and how governments are incorporating offensive cyber operations into their military strategy.

potential and actual attacks. Such sharing has proven difficult to establish for a range of reasons, including the fact that companies have been concerned about disclosing sensitive information, such as the financial data or communications of their clients. What is clear, however, is that the lines of communication within the private sector, as well as between the private sector and the public sector, require improvement.

Incentivizing Behavioral Changes: Multiple participants highlighted what one called "the inability to get the basics done" in the realm of cybersecurity. We know basic steps can be taken at the individual level, and through to companies and governments, to dramatically reduce vulnerabilities. And yet, even for something as simple as password protection, we continue to fail to take those fundamental steps. By way of illustration, one participant noted that while it had been 30 years since the first major computer worm (the Morris worm) was unleashed through the Internet and wrought havoc, the vulnerabilities revealed in that original attack are similar to many of those exploited in attacks today. Technological solutions are often easier to devise, participants noted, than ways of effectively changing human behavior, or what cyber experts often refer to as better cyber "hygiene." As one participant noted, we are not investing the needed resources, nor providing sufficient incentives, to address these vulnerabilities.

In considering incentives for adjusting human behavior, participants noted that it is important to take into account the different capacities of public and private sector actors. For example, a small business or a county election office is unlikely to have the expertise or resources to prevent or even detect an attack launched by a group of hackers or a foreign government, yet these smaller entities receive little support to defend against or respond to such outsized threats. When it comes to big companies, however, a participant noted that the risk associated with cyber attacks can be substantially mitigated if the necessary resources are expended. Nevertheless, most organizations are not willing to invest what is required, the participant said, and as it stands now, there appear to be insufficient incentives for large companies that have the knowledge and resources to consistently make the required investment across critical industry sectors.

When discussing regulatory options and other forms of interventions intended to promote cybersecurity through the use of incentives, participants spoke of the need to strike a balance between tapping the Internet's potential to power economic growth, and minimizing the inherent security risks. In essence, that means ensuring that companies can innovate and get products to market with reasonable efficiency, while at the same time appropriately prioritizing security and privacy. Participants noted a similar tension with respect to content moderation by social media platforms: while insufficient moderation can allow malicious actors to deepen cleavages, spread disinformation, and undermine the social fabric of our societies, so too can overregulation stifle free speech. One participant observed that every new communications technology has brought with it the fear that the medium would be used by nefarious actors to manipulate uneducated segments of the population, and in each instance prior to the advent of the Internet, consumers have proven to be savvier than anticipated.

Developing Norms and a Strategy for Deterrence: Participants noted that as an increasing number of states and non-state actors develop the operational capability to pursue their objectives through cyberspace – including the pursuit of malicious cyber activities that have impact in other states' territories –norms and a complementary strategy for deterrence are essential to prevent cyberspace from becoming a source of instability that could lead to inter-state conflict. Developing norms, supported by a strategy to deter actors from violating those norms, would assist in reducing the risks of misperception, miscalculations, and escalation; deter aggressive action; promote predictability and stability; and ultimately foster collaboration among international cyber actors to reduce our collective vulnerability in a way that enhances stability. Nevertheless, developing and promoting norms and deterrence – while preserving the benefits of connectivity – has proven to be quite challenging. Among the obstacles are: the wide range of stakeholders in cyberspace; the existence of rapidly developing technologies that continue to change the options for managing and framing cyber norms; the lack of traditional jurisdictional markers within cyberspace; the ease with which private and public actors can take action and have impact across borders; and concerns among major stakeholders about exposing their own sources, methods, and capabilities.

II. Working Group Discussions

Three of the Forum working groups were focused on the key vulnerabilities raised: critical infrastructure, privacy, and information challenges in social media. Additionally, we focused two working groups on areas in which we might promote policy designed to influence the behavior of responsible users – from governments to businesses to individuals – in an effort to better detect, deter, prevent, disrupt, degrade, and respond to the efforts of malicious actors. These included establishing and incentivizing the application of standards, benchmarks, and best practices; the development of norms and deterrence for public and private actors; and the education of individuals regarding the importance of exercising good cyber hygiene and risk management in business processes.

Each working group of about seven or eight participants evaluated at least two projects that had been developed by participants, in collaboration with CWP staff, prior to the Forum. Every project was developed with the idea of: (i) bringing multidisciplinary academic research and scholarship to bear on an aspect of cybersecurity in some significant way; (ii) in partnership with a practitioner (from outside of academia); (iii) in an effort to produce measurable impact within roughly three to five years; and all the while (iv) enriching research and scholarship. The lead drafter of each proposal was asked to present a succinct summary of the project, after which the moderator facilitated a discussion to provide critical feedback. Of particular importance were the following questions:

- Strengths and weaknesses. In particular, are there key weaknesses, omissions, or risks in the framing of the problem or the proposed solution?

- Implementation challenges. What are the greatest obstacles to effectively implementing this project, and can they be overcome?
- Potential improvements. How can the project be strengthened and its chances of success increased?
- Likely impact. If successful, what magnitude of impact will the project likely have on cybersecurity? Is the project scalable?
- Role of the university. Is research and scholarship important to the success of the project?

Participants were encouraged to think about how the projects they were considering could address issues being taken up by other groups beyond their own. For example, participants discussing a proposal intended to improve the cybersecurity of critical infrastructure were also asked to consider what implications it might have for privacy.

Before reconvening with the other Forum participants, the working groups were asked to prioritize the projects they reviewed, from the perspective of which projects CWP should pursue, and summarize the main points and any recommendations they wished to make regarding each project.

1. Critical Infrastructure

Critical infrastructure, the physical and virtual infrastructure that underpins essential services in our society, has become increasingly reliant on connectivity and digital technology, and consequently is increasingly vulnerable to malicious cyber activities. Over the last several years, we have seen evidence of this in the form of reported cyber attacks on and intrusions into infrastructure that continue to increase in number and scope. Among the most prominent have been the 2015 attack on Ukraine's power grid, which left 230,000 people in the West of the country without power for hours; the 2015 and 2016 thefts from the SWIFT global messaging system, a network employed by financial institutions to move money around the world; and the so-called 2017 WannaCry attack that temporarily paralyzed the British health system, among other affected entities – to cite just a few activities conducted by state and non-state actors. Perhaps the most powerful examples of recent intrusions into critical infrastructure can be found in the context of Russian efforts to interfere in the 2016 U.S. presidential election, which was in part focused on undermining the credibility of the U.S. election process; a June 2017 cyberattack delivered through Ukrainian accounting software, which wiped out data from the country's hospitals, airports, energy companies, banks, and federal agencies, as well as data critical to the operations of a Danish shipping conglomerate, an American pharmaceutical giant, and French manufacturing company; and the March 2018 report from the U.S. Department of Homeland Security indicating that Russian hackers had accessed machines

related to critical control systems at power plants in the United States.[9] The U.S. Department of Defense's (DoD) most recent Cyber Strategy, in fact, assumes that "a potential adversary will seek to target U.S. or allied critical infrastructure and military networks to gain a strategic advantage" during a conflict.[10] Nevertheless, in 2017, DoD's Task Force on Cyber Deterrence concluded that, "the unfortunate reality is that, for at least the next decade, the offensive cyber capabilities of our most capable adversaries are likely to far exceed the United States' ability to defend key critical infrastructures."[11] In other words, not only is our critical infrastructure becoming more vulnerable, and are intrusions increasing in scope and effect, but the sophistication of the malicious actors in this space is increasing as well. As a result, improving the cybersecurity of critical infrastructure is an urgent priority.

In the critical infrastructure working group, four projects were discussed. The first was focused on election infrastructure; the second on reducing the vulnerability of digital microelectronics to supply chain tampering; the third on information sharing among key critical infrastructure partners across the private and public sector; and the fourth on reducing the vulnerability to cyber attacks of computers that interact with the physical world, such as utility infrastructure, modes of transportation, and medical devices. Each project is summarized below, along with a few of the issues raised during the course of the discussion in the working group.

A Robust Election Auditing System: This project idea is aimed at improving the security of elections while also improving public confidence in the broader integrity of those elections. The proposal is to pilot a new auditing approach in one or more jurisdictions in the United States that would assist in identifying vulnerabilities in any election process, increase voters' confidence in the election results, and drive the development of an election process that cannot be subverted without detection.

The auditing approach would combine two methodologies of auditing. The first would be an administrative "risk-limiting audit" (RLA) that typically involves election workers randomly sampling physical ballots and manually confirming that their contents are consistent with expectations given the results of the election. The second, called "end-to-end verifiability" (E2E verifiability), would allow voters to verify after an election that their votes had been counted and recorded properly, without disclosing to anyone how they voted. The pilot would be monitored, and an assessment done as to the effectiveness of the audit, its use, and the impact the process has on voters and election administrators. If successful, this two-pronged approach to auditing elections might be replicated in other jurisdictions.

---

[9] U.S. Department of Homeland Security, "Russian Government Cyber Activity Targeting Energy and Other Critical Infrastructure Sectors," March 16, 2018, https://www.us-cert.gov/ncas/alerts/TA18-074A.

[10] U.S. Department of Defense, *The DOD Cyber Strategy*, April 2015, http://archive.defense.gov/home/features/2015/0415_cyber-strategy/final_2015_dod_cyber_strategy_for_web.pdf.

[11] U.S. Department of Defense, Defense Science Board, *Task Force on Cyber Deterrence*, February 2017, https://www.acq.osd.mil/dsb/reports/2010s/DSB-cyberDeterrenceReport_02-28-17_Final.pdf.

Key questions and concerns raised by participants included: (1) the cost of establishing the proposed auditing system, and whether most jurisdictions would be able to afford it; (2) whether the technical understanding of how to employ E2E-verifiable methodology is sufficiently advanced so as to allow for it to be deployed within a three to five-year timeframe, which is a requirement for CWP; and (3) whether election administrators would be open to trying this approach, which, in addition to being complex, would expose their systems to outside review in a way that could expose other problems (including ones not caused by attempted election meddling). On the issue of cost, participants discussed whether a prior attempt to implement a multi-pronged auditing approach in Travis County, Texas, might offer important lessons learned, and furthermore, whether a pilot might be scaled in such a way as to reduce the costs of this approach when adopted by multiple jurisdictions and vendors.

Reducing Vulnerabilities in the Supply Chain: This project would attempt to address a key vulnerability for critical infrastructure, which is heavily dependent on digital microelectronics that often have a complex supply chain involving multiple industries, manufacturers, and distributors in a wide range of countries. Adding to the vulnerability, the nature of hardware design encourages the reuse of microelectronics, which means tampering with one design could provide a foothold for attacking millions of devices. This project seeks to mitigate some of the risks inherent in a supply chain that is built on untrustworthy parts and untrustworthy personnel by applying new techniques that make digital microelectronics more resistant to tampering and that test their trustworthiness. The techniques developed by a Columbia University professor and his students include algorithms that can help detect vulnerabilities in the supply chain and crafting designs that make it more challenging to tamper with the hardware. Participants agreed that this is a significant vulnerability that needs to be addressed and that it would be ideal if these techniques could be piloted with both a public and a private sector partner.

Real-Time Information Sharing on Critical Infrastructure Threats: In an increasingly complex and dynamic cyber threat landscape, one of the most important capabilities for protecting our critical infrastructure is the ability to share information among critical infrastructure partners in the private sector and relevant government entities. The third project proposal involves piloting a system that would facilitate real-time, automated sharing of relevant, timely, and actionable threat information across the financial sector on cyber fraud – not only to enable proactive responses to attacks when they occur, but also to provide predictive analysis that may help prevent such attacks. In particular, the project would initially seek to enable information sharing across financial institutions about low or no-volume accounts that are opened and kept dormant, which are often used to commit fraud. The project would do this using a technology developed by a professor at Columbia University that would employ Bloom filters to facilitate the sharing of account information among participating financial institutions *without* revealing the identity of institutions' clients, which is a concern for banks. The project would also look to companies to help identify additional corroborating indicators that are effective in spotting and

predicting fraud, and find ways to share information about these indicators by integrating with, and building on, existing fraud detection systems. If effectively piloted, this approach could be scaled across other financial cyber threat scenarios, as well as to information sharing in other critical infrastructure areas. Ideally, the project would engage both industry and government partners, including the Financial Systemic Analysis and Resilience Center (FSARC), which could help provide avenues of secure, automated communication. Participants agreed that the current level of sharing of information on cyber threats is inadequate, and one participant noted that automatic sharing of a discrete and previously agreed to set of information could help to compensate for institutions' general inclination not to share.

Securing Cyber-Physical Systems (CPSs): The fourth and final project seeks to address a vulnerability of legacy CPSs, or computers that interact with the physical world, which may be exploited to inflict significant harm. The proposed project would involve deploying in various scenarios a technology developed by a professor at Columbia University to reboot CPSs approximately once every few seconds (or even every few microseconds), to clear tainted state or bad inputs provided by a cyber attacker, and to thereby limit the windows within which an attack may be carried out. While new CPSs could be designed to be entirely stateless, some existing stateful systems can be better secured by repeatedly rebooting to clear state. The reboots are so swift – taking just one quarter of a millisecond to complete in the case of a drone, for example – that the systems can continue to function during that time on the basis of inertia. And the period between reboots – occurring approximately once every few seconds – is short enough to make it much more difficult for attackers to gain a foothold in the system's memory. Also critical is the fact that each reboot creates a fresh copy of the software to start again. The replica function means that the operating program is always changing, consistently forcing the attacker to start from scratch and find new vulnerabilities. Participants explored possible use cases for this project and asked why this was preferred to other options, such as using other firmware security approaches.

2. Privacy

As more of our daily lives are connected to the digital world, protecting our personal information online has become increasingly challenging. Recent high-profile data breaches conducted by state and non-state actors demonstrate how vulnerable many entities are that hold sensitive personal data,[12] though such breaches are only part of the problem. A significant amount of personal information about individuals can be obtained through open sources, or

[12] Obvious examples include the U.S. Office of Personnel Management data breach in 2014, in which the private information of more than 22 million current and former U.S. federal government employees, contractors, and their friends and families was compromised, as well as the 2015 breach associated with a series of Canadian government websites, an attack claimed by hackers associated with Anonymous. "Government of Canada websites under attack, hacker group Anonymous claims responsibility," *National Post*, June 17, 2015, https://nationalpost.com/news/canada/government-of-canada-websites-under-attack-environment-canada-foreign-affairs-down.

with the willful participation of firms and platforms that gather such information. Particularly when such data is aggregated, significant privacy issues arise that are unique to the digital age.[13] In addition to the legal, values-based, and ethical reasons to want to better protect the privacy of individuals, there are also economic and social incentives: if people lose trust in the capacity and will of public and private entities to protect their information and use it for legitimate purposes, an increasing number of people may choose to reduce their use of the Internet, posing a serious threat to economic growth and other avenues of empowerment it enables.[14] Of course, the two projects considered by this working group – as well as the information-sharing project presented in the critical infrastructure working group – demonstrate how new and evolving technologies can also *promote* privacy. Yet, as one participant observed, the development of such affirmative applications has up to this point been relatively limited and merits greater investment.

Before reviewing the two projects assigned to this group, participants discussed several conceptual and definitional issues around online privacy, beginning by distinguishing between personally identifiable information, which is information that could potentially identify a specific individual, and the information individuals generate – knowingly and unknowingly – through online behavior. Participants discussed what people's expectations are with respect to these different categories of information and the extent to which governments, companies, and individuals should be responsible for protecting such information. Then, the group turned to the two projects under consideration.

Giving Online Users Greater Control Over Their Data: The first project was designed to assist in shifting control over the data that users generate on the Internet through online browsing, for example, back into the hands of individuals, using a technology called MixIt. Developed by

[13] Several academics have noted that due to technological developments over the last several decades, entities have become much more skilled in recording data that is produced through our daily interactions with various institutions, such as toll booths, grocery stores, and drugstore chains. Although such information may have been collected to provide individual users with better service, much of the time those users are not provided with access to the information collected about them, a phenomenon academics have termed "inverse privacy." See, e.g., Yuri Gurevich, Efim Hudis, and Jeannette Wing, "Inverse Privacy," Communications of the ACM, July 2016, Vol. 59 No. 7, https://cacm.acm.org/magazines/2016/7/204020-inverse-privacy/abstract. This practice raises at least two major concerns: (i) third parties are increasingly collecting and retaining information regarding our daily lives that, in aggregate, could raise significant privacy concerns; and (ii) individuals whose data is being collected might be able to use that information for other purposes. The latter concern raises questions as to whether individuals under such circumstances should have access to the information collected about them, and should be informed about, and perhaps even exercise some control over, its intended use and retention.

[14] Evidence of such a decline in trust abounds. For example, a recent survey of consumer perspectives on cybersecurity and associated privacy risks, found that consumer trust in businesses to secure their personal information is fading, with only 25 percent of respondents believing that most companies handle their sensitive personal data responsibly, and 85 percent of consumers indicating that they will not do business with a company if they have concerns about its security practices. Meanwhile, government actors enjoy even less trust – as 72 percent of individuals believe companies are better equipped than government to protect their data. See, for example, PricewaterhouseCoopers, "Consumer Intelligence Series: Protect.me," September 2017, https://www.pwc.com/us/en/advisory-services/publications/consumer-intelligence-series/protect-me/cis-protect-me-findings.pdf.

a Columbia University professor and his students, MixIt operates continuously in a user's browser to create "cover traffic," making it nearly impossible for the user's real browsing data to be discerned and collected without her permission. As a consequence, this software could – if deployed widely enough -- shift the current balance of power over a user's browser activity on the Internet and allow the user to have direct control over the information they produce, as well as perhaps enable users to generate revenue by selling such data to companies that wish to use it for commercial purposes. Participants discussed how many users would be necessary to make MixIt technology work, and whether a revenue generating model could be established through which thousands of users, for example, would be capable of banding together to sell data regarding their activity on the Internet, while recognizing that, if effective, the model would pose a challenge to the dominant business model of online platforms and some other companies.

Key issues and concerns raised by participants included: (i) the impact on the economics of the current Internet service providers, and whether generating the critical mass of MixIt users would generate an unintended cost for users and non-users (e.g., reducing speed of access); (ii) how an individual's unique browsing data could be catalogued and eventually shared in a way that would be economically useful; (iii) how one would best develop a market for such data; (iv) whether Internet service providers might seek to charge individuals who use MixIt (given the additional bandwidth use), or if the platforms and other companies that currently offer online services for free – through monetizing the very data that MixIt would aim to prevent them from collecting – would need to pass along the cost to consumers; and (5) whether the success of this project would mean a loss of big data that is currently useful to the public.

Building a Digital Identity System: The second project proposed bringing together an interdisciplinary team to design and pilot an efficient, secure, decentralized digital identity system in an area where a significant number of people do not have access to official identifications. A trusted and verifiable digital identity is essential in modern digital societies and economies, and can serve as an official identification, which is often fundamental to a government's ability to deliver vital services to its people, and in many societies can affect an individual's ability to do everything from visiting a doctor, to voting in an election, to enrolling in school. It is also critical to accessing the private sector – financial institutions often require an ID to give loans and credit, and to allow people to open a bank account. Yet more than a billion people worldwide do not have access to an official identification, which in many cases results in their not having access to resources to which they are entitled. The project envisions working with a range of partners – including national government, industry, citizen groups, and experts on the local culture and the different aspects of technology (particularly smart phones, cyber security, and cryptography) – to build and pilot a system that would meet the following criteria:

- universal (available to all without undue hardship involved in acquiring the capability and usable among all participating entities without additional identity verification);

- useful (fill a critical void, provide a significant enhancement to existing capabilities, or create opportunities for new irresistible applications to justify the cost);
- usable (easy to use, maintain, and guide users through a transparent process in which she decides how much or how little information to reveal for a particular transaction);
- secure (from fraud and abuse by criminals, corporations, or the national government; in addition, a lost device must be unusable by anyone other than the original owner, and must be straightforward to replace);
- privacy-enhancing (should not require a centralized database that could be used to track user activity or commit fraud; should not force consumers to surrender anonymity where that would be inappropriate; and should give users the right not to use the system and to access an alternative ID system); and
- affordable (should leverage existing technology platforms, networks, physical facilities, and institutions to help reduce costs).

Any identity system would have to overcome well-justified societal concerns about individuals providing governments or corporations with too much visibility into their private lives. Other concerns include the fear of vulnerabilities that could result in the irretrievable loss of control over one's identity, and significant implementation costs and technical challenges, particularly for any national-scale deployment and management of public key infrastructure.

In discussing the proposal, participants noted that it would be important to learn from and build upon the systems of digital identity have been deployed in other countries, and in particular, from the problems that have surfaced in their implementation (the challenges with India's biometric system being one example). Another participant noted that the system could provide a useful way to cut down on the inconvenience of managing multiple forms of identification and verification, and render existing indices that contain personally identifiable information less vulnerable to misuse. Several participants said one of the proposal's chief strengths was developing the capacity for different kinds of transactions to be carried out with different tiers of personal information, so that the only disclosure made in the context of every transaction is what is absolutely necessary and no more. Participants noted that while we have much of the technology we need to build a secure, decentralized virtual ID, a major challenge would be scaling the public key infrastructures needed to apply such technology, and building sufficient public trust in such a system. Another participant made the case that it would be important not to consolidate virtual ID information in a single place, such as with a company or a database, so as to reduce the risk of it being compromised. Lastly, a participant recommended that, in addition to developing the technology for a virtual ID and conducting the behavioral, cultural, and public policy research to improve the likelihood of its uptake, it would also be important for such a project to develop a set of norms and principles, which could inform regulatory frameworks regarding the use of personal data and identity systems. Without such an effort, the participant said, public or private actors could develop systems (independent of

the proposed digital ID system) that could identify and even track individuals without their knowledge or consent, among other forms of misuse.

3. Information Challenges in Social Media

As is often true in the realm of the Internet, one of its greatest assets – the ability to provide almost immediate access to immense troves of information – has also become one of its greatest liabilities. Efforts to manipulate or fabricate information, of course, long predate the advent of the Internet. But the way information is shared and consumed on the Internet, together with rapid advances in technology, has made the medium particularly ripe for efforts to undermine the integrity of content, as well as to track, and abuse, information regarding the behavior of individuals. This in turn has the capacity to undermine not only the utility of the Internet, but also to exacerbate divisions in our societies and to weaken liberal democracies by degrading, among other things, our capacity for having a productive, fact-based public dialogue regarding issues of consequence.

This is especially true in light of the growing proportion of people around the world who use social media platforms to get their news and other information about the issues that matter most to them[15] and the fact that the spread of misinformation and disinformation on social media is increasing, at least in part because technology has made it easier to create and rapidly disseminate misinformation and disinformation, and harder for users – and in some instances even experts – to identify such false content. Indeed, the reach and penetration of these social media platforms (the largest, Facebook, has 2.23 billion users),[16] the frequency with which users visit them (approximately three-quarters of Americans visit Facebook every day),[17] and the amount of time users spend on them (the average U.S. user spends more than 50 minutes on Facebook every day),[18] has made such platforms one of the most important – if not the predominant – virtual public square, and one whose reach extends across national borders. Consequently, how these platforms address emerging dilemmas around information integrity has a profound impact on everything from public discourse and user knowledge, to the way we understand fundamental issues like free speech.

---

[15] According to a Pew Research Center Survey, approximately two-thirds of Americans adults say they at least occasionally get their news on social media. Elisa Shearer and Katerina Eva Matsa, "News Use Across Social Media Platforms 2018," Pew Research Center, September 10, 2018, http://www.journalism.org/2018/09/10/news-use-across-social-media-platforms-2018/. Approximately three-quarters of Americans visit Facebook every day. Pew Research Center, *Social Media Use in 2018*, March 1, 2018, http://www.pewinternet.org/2018/03/01/social-media-use-in-2018/.

[16] "Company Info," Facebook, accessed September 18, 2018, https://newsroom.fb.com/company-info/.

[17] Pew Research Center, *Social Media Use in 2018*, March 1, 2018, http://www.pewinternet.org/2018/03/01/social-media-use-in-2018/.

[18] James B. Stewart, "Facebook Has 50 Minutes of Your Time Each Day. It Wants More." *The New York Times*, May 5, 2016, https://www.nytimes.com/2016/05/06/business/facebook-bends-the-rules-of-audience-engagement-to-its-advantage.html.

The concerns associated with misinformation and disinformation are exacerbated by other issues that arise in social media environments, which cannot easily be disaggregated if we are to identify sustainable and effective responses. For example, as online social networks play a growing role in shaping not only users' views but also their relationships and communities, there are growing concerns that the way social media platforms present content to users, and the algorithms they often use to order such content, are further aggravating existing cleavages in society – presenting users with filtered content that affirms their own views, and even perhaps driving users towards more extreme views and content. This in turn may undermine some of the critical elements of liberal democracies, such as the ability to agree on shared arbiters of basic facts, and the exposure to ideas that are different from one's own. Another central element to understanding content accurately comes from relevant context that is frequently not explicit in an online environment – such as, who is paying for a message to be disseminated to a given user, and why has that individual been targeted as a recipient? Such context can help the user judge, among other things, the credibility of a message, particularly when it appears in the context of political campaigns.

Furthermore, while social media platforms in many ways function as a virtual public square, the human and algorithmic decisions that are made regarding their operation are not fully transparent – information that may well be critical to understanding, and potentially improving upon, the way they function. Of course, sharing such information would need to take into account the privacy of users, but greater transparency regarding how these platforms work is essential to understanding how they are shaping our societies and individual behavior. The working group on information challenges considered two project proposals intended to assist in addressing the foregoing problems.

Social Media Solutions Lab: This project proposal envisioned two methods for developing and testing potential solutions to specific information challenges in social media identified above: (i) working inside existing social media platforms; and (ii) working outside of existing social media platforms. For the inside approach, the concept would be to work with one or several of the big platforms to develop and test solutions for key information challenges, such as ways to mitigate the spread of misinformation and disinformation; to increase the credibility of accurate information; or to enable users to select different ways of ordering, labeling, or filtering the content that appears in their feeds – distinct from the order prescribed by the platform's algorithm. The feedback from users testing out such options, or "filters," could then be compared to those not using such options or the filters, in an effort to gauge whether any particular tool is effective. For the outside approach, a partner platform serving a bounded community would offer an "open protocol" system, whereby users and third parties would be able to design different filters that would give individuals the power to curate their own feeds, or prioritize the way content is ordered according to different metrics (such as giving preference to accuracy or a diversity of views). This approach would effectively push control and power out to the ends of the network, rather than monopolizing such decisions in the

center, as is the default with the big platforms. In addition, the lab would allow third parties to compete to provide better overall interfaces for viewing the content and incentivize the continual improvement of such solutions, in order to try to attract more users.

Key issues raised by participants included: (i) whether, given the choice, users would want to see a broad spectrum of viewpoints, or would prefer to see information and views that support their own; (ii) similarly, whether users would prioritize accuracy when they consume content on social media; (iii) whether the "inside" approach was viable, given that it would likely cut against the existing platforms' business models (even though a decentralized model might be advantageous to platforms, given that handing over more decision-making power to individual users would allow platforms to be perceived as less of an arbiter); (iv) what other solutions for information challenges might be tested in such a lab scenario; and (v) for the "outside" approach, whether the ideas discussed, including new methods of filtering information, could be adequately tested without the collaboration of the big platforms (in part because researching these ideas at scale is critical to both understanding the problem and testing solutions, given that some information challenges only emerge among a user base that is larger than the one likely to be attracted to a test bed).

A Platform for Addressing Information Challenges: The second project would aim to help individual users process, filter, and understand information they consume online, by giving social media users the ability to share – in a secure, anonymous way – information that they generate through their use of Internet. The information shared by individuals would be aggregated on a platform in a way that is accessible, user-friendly, and transparent, where it could then be used by researchers and developers to build tools aimed at addressing the most common challenges revealed in the shared data, much of which is currently only in the hands of social media companies. Tools developed on the platform would, in turn, be offered to the full spectrum of social media users, to layer atop their platforms in a way that would improve their information consumption. Issues and concerns raised by participants included: (i) whether it would be possible to convince existing social media platforms to share the necessary data regarding users' behavior (with users' permission, of course), develop a rigorous methodology to "read" an individual's information consumption within existing data structures for these purposes, and then make the information accessible to those wishing to develop tools for users; (ii) whether it would be possible to build exemplar tools to tackle known issues (such as revealing deep fakes and flagging content generated by foreign influence campaigns); and (iii) the importance of agreeing on a framework for measuring whether the tools are working.

Ultimately, the group settled on a hybrid of the proposals, which had two parts. First, creating a solutions laboratory that would bring together an interdisciplinary mix of software engineers, policy experts, social and behavioral scientists, journalists, and others to develop a set of applications aimed at giving users ways to improve their consumption of information and remedy crucial challenges such as disinformation campaigns; piloting these applications, and conducting research on their efficacy; and providing a forum for reaching out to the social

media companies to press for the adoption of tools proven to be effective. And second, working with one of the major platforms to develop a set of filters, which would be tested with a subset of willing users, in order to measure the impact of giving individuals greater power in shaping their social media feeds. Ideally, those filters that proved useful could then be offered to the platform's full network of users. Among the platforms that might be approached, there was a strong view that Twitter should be considered in light of its previous efforts to tackle some of these problems; previous applications that have been created to filter content on Twitter; and the openness the company has demonstrated its approach and work with others on the presentation of information. Academics also noted that regardless of the partner, it would be important to negotiate up front the extent to which data received from a platform can be shared publicly, so that researchers are able to publish findings.

4. Standards, Benchmarks, and Best Practices

As was discussed in the initial plenary by participants, one of the most significant steps that could be taken to improve cybersecurity across the board would be for organizations and individuals to prioritize basic cybersecurity practices. The current failure to do so is widespread, well documented, and generally decried, yet no clear path to solving the problem has been identified. As such, one of the main questions raised was whether an appropriate set of incentives could be established to promote the adoption of standards, benchmarks, and best practices when it comes to cybersecurity. [19]

The working group discussed three projects aimed at motivating private and public sector actors and individuals to prioritize good cybersecurity practices and better manage the associated risks. At the outset, the working group identified several major barriers to the development, adoption, and application of effective cybersecurity standards, including: (i) the monetary and human resource costs associated with compliance; (ii) the fact that individual users – who have a critical role to play in prevention, detection, and resilience – often lack the

[19] The International Organization for Standardization (ISO) – a nongovernmental organization that develops voluntary, consensus-based, and market relevant international standards – defines a standard as "a document, established by consensus and approved by a recognized body, that provides, for common and repeated use, rules, guidelines, or characteristics for activities or their results, aimed at the achievement of the optimum degree of order in a given context." Although further work on standards that meet this definition would be useful, numerous standards have already been developed for cybersecurity to help organizations better manage security risk and implement security controls. The cybersecurity framework created by the National Institute of Standards and Technology (NIST) is one example of a framework that sets out a core set of activities that organizations of all kinds and sizes can take to assess their cyber risks and take action to reduce them, focusing on five key functions: identify, protect, detect, respond, and recover. Private sector actors have also formed their own structures and mechanisms to share information in order to reduce risk and increase resilience – and to coordinate cooperation with the government – such as the Information Sharing and Analysis Centers, or ISACs, which are organized by sector (e.g. financial services and water). See "Standards in Our World," International Organization for Standardization, accessed September 18, 2018, https://www.iso.org/sites/ConsumersStandards/1_standards.html; National Institute of Standards and Technology, "Cybersecurity Framework," U.S. Department of Commerce, accessed September 18, 2018, https://www.nist.gov/cyberframework.

tools, knowledge, and incentives to play that role; (iii) a crowded market, with various frameworks for incentivizing cybersecurity practices, but no one-size-fits-all solution;[20] (iv) the difficulty of effectively implementing better cybersecurity practices, whether because of the continuing reliance on legacy systems,[21] longstanding business practices that inhibit adoption, or other reasons; and (v) the difficulty of measuring the efficacy of, and compliance with, such frameworks. The working group additionally discussed the particular need to focus further work on establishing cybersecurity standards and practices for the development of devices that are part of the Internet of Things, which are becoming more ubiquitous by the day. Experts predict more than half of all new businesses will make use of the Internet of Things by 2020,[22] at which time it is estimated some 20 billion such devices could be deployed worldwide.[23] Yet the lack of adequate measures to secure the Internet of Things is making it easier for hackers to compromise the growing number of connected devices through the weakest link. Established best practices for the Internet of Things might be tied to a similar incentive structure as the one applied to organizations and individual users, or it could be different (for example, relying on a labelling or rating system managed by a third-party assessor that consumers would find trustworthy).

A Cybersecurity Bond Rating: The first project proposal focused on incentivizing the private and public sectors to improve their cybersecurity practices by promoting the further development of quantifiable Benchmarks, Standards, Guidelines, and Best Practices (BSGBs), including for the Internet of Things. The project would then pilot the promotion of these BSGBs through a rating system that would evaluate entities' adoption of such standards by virtue of an auditing process. Measurable indicators would be developed to assist in the auditing process, and the BSGBs would be published to ensure consistency and signal to the market what was being considered. Rating agencies would evaluate the "cyber-worthiness" of a product or entity based on their assessed adoption of BSGBs, either in the context of an overall rating or in a separate bond product, as has been done in the context of "green bonds," which

[20] The National Institute of Standards and Technology (NIST) framework being a key exception, as a widely used framework that sets out a core set of activities that organizations of all kinds and sizes can take to assess their cyber risks and act to reduce them.

[21] As one expert noted, even if an entity chooses to implement controls based on a framework like NIST, implementing those controls across all of the technology upon which that entity relies is often complicated by the use of outdated "legacy systems." This technology tends to be less secure because it is no longer subject to regular vulnerability patching and can be incompatible with controls needed to defend against evolving threats. (Entities tend to keep such systems because they are viewed as too critical to take down for even a short period of time, or because upkeep is seen as less costly than investing in new systems.) As a result, entities may implement controls on some but not all of their systems, potentially undermining the efficacy of adopting and implementing a framework like NIST.

[22] Gartner, "Gartner Says By 2020, More Than Half of Major New Business Processes and Systems Will Incorporate Some Element of the Internet of Things," January 14, 2016, https://www.gartner.com/newsroom/id/3185623.

[23] Liam Tung, "IoT Devices Will Outnumber the World's Population This Year for the First Time," ZDNet.com, February 7, 2017, https://www.zdnet.com/article/iot-devices-will-outnumber-the-worlds-population-this-year-for-the-first-time/.

attempt to evaluate the environmental impact of certain bond-issuers. In theory, companies desirous of a good cyber-worthiness rating would adopt the relevant BSGBs, thereby improving their, and their customers, cybersecurity.

Key issues and concerns raised by participants included: (i) whether this structure would effectively incentivize entities to fully implement BSGBs; (ii) whether a financial services framework would be capable of collecting and integrating relevant data into a pricing model within the project's desired time horizon of three to five years; (iii) whether it would be possible to establish a sufficiently flexible set of BSGBs to reflect the range of actors subject to such criteria; (iv) whether the BSGBs could be both effective and sufficiently affordable for smaller public and private entities to adopt; (v) what entity should bear primary responsibility for managing the rating system; and (vi) whether it would be possible to induce private companies to disclose sufficient information regarding cyber incidents and practices to conduct a credible audit. Overall, participants indicated that flexible criteria capable of taking into account various entities' different vulnerabilities, resources, and risk tolerances would be important; that the use of relative metrics that would allow entities to measure themselves against similarly situated organizations would be useful; that the auditing process should be automated to the greatest extent possible, in order to reduce the cost for participating entities; and that a collaboration between an academic institution, such as Columbia University, and a traditional rating agency could be beneficial.

Cybersecurity Insurance: The second project was complementary to, but distinct from, the first in that it proposed developing uniform metrics and an auditing process for measuring adherence to defined BSGBs, in order to provide the foundation for a cybersecurity insurance market. The institution responsible for the audit would verify an entity's relative compliance with the BSGBs and certify the entity before an insurance underwriter. Insurance agencies would then be able to offer cybersecurity insurance at a premium rate for approved entities. With Allianz, a leading integrated financial services provider, predicting that the global premiums for cybersecurity insurance are expected to exceed $20 billion by 2025,[24] the ability to reduce premiums would provide an incentive for companies that wish to purchase cybersecurity insurance to improve their cybersecurity performance, while additionally lowering the risk and vetting costs for insurance providers. That said, participants questioned what prior evidence of risk reduction would be required for the insurance market to lower costs for entities that implement certain cybersecurity measures.

Cybersecurity Education for Children: The third project would seek to work with an educational partner to design and pilot a comprehensive cybersecurity curriculum aimed at enhancing the digital literacy of middle school students, with the goal of eventually rolling out the curriculum across a range of ages in public schools across a U.S. state. Children have unique

[24] Allianz SE, "Businesses must prepare for new generation of cyber risks," September 9, 2015, https://www.allianz.com/en_GB/press/news/studies/150909-businesses-must-prepare-for-cyber-risks.html.

vulnerabilities in cyberspace, which are only increasing as they use the Internet at younger ages, with greater frequency, and with less supervision. While repeated calls have been made for educating children (and their parents and teachers) about how to make better choices regarding their security, privacy, and health online, efforts to achieve this have so far been scattershot and ineffective – lacking rigor and breadth, and leaving schools with strained budgets and limited expertise to develop their own responses. The overarching goal of the proposed curriculum would be to give youth the knowledge and tools to use the Internet and technology safely – taking advantage of the opportunities, while also being aware of the risks; and to inculcate a culture of cybersecurity that promotes a more secure Internet. A potential partner would be Theorem Studios, which is developing a multi-platform storytelling experience focused on cybersecurity, including a series of 20 short films, Virtual Reality content, podcasting, and social media, which could be incorporated into the curriculum.

Participants queried how the impact of the curriculum could be tracked and measured on a scale that would be useful for research, so as to learn whether what was being taught translated into changes in behavior, and said more research was needed on whether such an effort is already being worked on by other institutions. Participants also highlighted several potential obstacles to – or complexities involved in – implementing this project, including: (i) the lack of instructors in middle and high schools qualified to teach such a curriculum; (ii) the decentralized nature of the U.S. educational system, which would preclude systematic dissemination of a curriculum; and (iii) the lack of consensus among policymakers, administrators, teachers, and even parents of the need for cybersecurity education. There was a consensus among the working group participants that the curriculum should not simply be online and that communicating the importance of a cybersecurity education to the full range of education and policy leaders would be vital. Independent of a CWP project, multiple participants suggested that Columbia could provide expertise for the Theorem Studios project aimed at improving public understanding of cybersecurity.

5. Norms and Deterrence

The norms and deterrence working group considered how the development of international norms and a strategy to promote deterrence could help promote security and stability in cyberspace. Although most countries have indicated that international law applies in this realm, considerable work remains to explain and to agree on how existing law and norms apply in the cyber sphere, both in peacetime and in the context of an armed conflict. Participants agreed that the development of additional cyber norms and activities designed to support such norms could help to reduce the risks of misperception, miscalculations, and escalation; deter aggressive action; and ultimately foster collaboration among cyber actors to reduce our collective vulnerability in a way that promotes predictability and enhances stability. The working group considered two projects in this area, with the first focused on building norms and the second on building a deterrence strategy.

Tech Accord Partnership: The first project would bring private companies together with academics to develop industry norms aimed at enhancing cybersecurity and ultimately promoting the development of norms for state actors. The proposed project would involve expanding and strengthening the Cybersecurity Tech Accord,[25] a public commitment by more than 60 global companies to protect users everywhere, oppose cyber attacks on innocent citizens and enterprises, and empower people to strengthen cybersecurity protection. The project would aim to improve and build upon of the norms promoted by the Tech Accord in two ways. First, by strengthening the existing norms – through analyzing ambiguities, exploring possibilities for refinement or elaboration, considering metrics for analyzing compliance, identifying and facilitating partnerships to promote industry norms currently endorsed by the Tech Accord, as well as identifying and sharing best practices associated with such norms. And second, by identifying areas for the development of additional norms, such as norms for the Internet of Things or for the operators of critical Internet infrastructure.

Participants underscored that – given the importance of the private sector to cybersecurity from a practical perspective in the digital realm, the power wielded by a handful of global companies that manage significant aspects of cyberspace, and the time required on average to create binding, global treaties – a private-sector-led initiative on cyber norms presents an appealing way to promote the development of effective norms for state-actors. Nevertheless, participants raised some questions regarding the nature of the proposed relationship between the academic group and the companies participating in the Tech Accord, and queried whether the power wielded collectively by the private sector could ever come close to that of government (and thus the degree to which any industry norms could have a significant impact on cybersecurity). These concerns notwithstanding, participants suggested several areas – beyond providing expert analysis on the technical and policy aspects of existing and new norms – where academic research and scholarship could be useful to the Tech Accord, including: (i) offering insights regarding effective industry self-regulatory efforts, drawing on prior historical examples; (ii) helping to develop some kind of independent oversight body, which would help to evaluate the success of the Tech Accord in affecting the behavior of industry entities and the promotion of broader norms; and (iii) creating some sort of accountability mechanism, perhaps simply through naming and shaming activity that is inconsistent with the Tech Accord.

Promoting a Strategy for Deterrence: The second project would bring together university researchers with partners such as the Cyber Threat Alliance, the National Intelligence Council, and the Cyber Threat Intelligence Integration Center to develop a standard and transparent methodology to help assess whether (and under what conditions) U.S. government deterrent

[25] The Cybersecurity Tech Accord, "About the Cybersecurity Tech Accord," https://cybertechaccord.org/about/. (Accessed November 27, 2018.) The four key principles adopted by signatories of the Tech Accord are: 1. We will protect all of our users and customers everywhere. 2. We will oppose cyberattacks on innocent citizens and enterprises from anywhere. 3. We will help empower users, customers and developers to strengthen cybersecurity protection. 4. We will partner with each other and with likeminded groups to enhance cybersecurity.

actions suppress future cyber attacks. This project would assist in evaluating the policy promoted in the new U.S. National Cyber Strategy[26] of imposing "swift, costly, and transparent consequences" when malicious actors harm the United States or its partners, which, it is presumed, will deter malicious actors from attacking U.S. assets. Participants pointed out that there is currently no methodology to measure whether such a policy will deter or provoke additional attacks, and thus a rigorous review of its impact will be critical to evaluating its success. This project would aim to develop and test such a methodology, updating it as the cyber landscape evolves, with the primary aim of facilitating unclassified assessments by the commercial cyber threat intelligence community, and the secondary aim of serving the U.S. Intelligence Community.

Working group members recognized the utility of formulating an objective standard of measurement to determine the effectiveness of different cyber deterrence strategies, as well as the value associated with collecting and analyzing data of this sort and making it accessible to other researchers. Nevertheless, participants raised concerns regarding the degree to which it would be possible to measure causality with respect to the impact of deterrence measures, given the multitude of variables that could lead to an acceleration or deceleration of cyber attacks by malicious actors, and given that the United States does not typically acknowledge offensive cyber activities intended to impose consequences on those who have attacked the United States. In response, one participant suggested that the methodology would most likely calculate ranges that, taken together, could at least suggest causal effects, which was a technique used by the Correlates of War project at the University of Michigan. Of course, that would not address the concern regarding whether any actions taken by the United States to impose consequences on malicious actors are likely to be disclosed by the United States or other governments in the context of a deterrence strategy – a vital data point for such an analysis. Another concern raised was whether this project would meet the criteria for CWP projects, as it would not clearly result in measurable impact unless the government decided to pursue a strategy or decision based on the research carried out (nor, for that matter, is it likely that a government would acknowledge using such research, even if it were to do so).

III. Conclusions and Project Selection

When participants reconvened in a plenary session, the five moderators reported out on the project ideas discussed in their respective working groups, describing the cybersecurity challenge each project would seek to address and the working group's assessment of its main strengths and weaknesses. Participants were then given a chance to ask follow-up questions about projects that had been discussed in other working groups.

[26] The President of the United States, *National Cyber Strategy of the United States of America*, September 2018, https://www.whitehouse.gov/wp-content/uploads/2018/09/National-Cyber-Strategy.pdf.

Next, each participant was asked to identify the two or three projects that she or he thought were most promising out of all of those described, and to provide recommendations regarding which projects should be further developed by CWP for possible funding and implementation. In weighing the projects that most merited pursuing, participants were reminded of CWP's key criteria: namely, that projects should bring multidisciplinary academic research and scholarship to bear on the challenge of cybersecurity, in partnership with a non-academic entity, in an effort to produce measurable impact within roughly three to five years – while simultaneously enriching research and scholarship.

It was notable that virtually every project received support from one or more participants, which speaks to the quality of ideas proposed and the breadth of cybersecurity challenges worthy of attention. Nevertheless, three projects proposals received the greatest support for further development by CWP: (i) a digital identity system, (ii) a robust election auditing system, and (iii) an online laboratory for testing solutions to the information and privacy challenges in social media and other online activity. In addition, several other project ideas were identified as having significant potential, whose further development CWP might encourage other parts of Columbia University and outside partners to carry forward.

The first project would bring together an interdisciplinary team to design and pilot an efficient, secure, decentralized digital identity system that could be deployed in a country where a significant number of people do not have access to official identification, and as a consequence are unable to access resources to which they are entitled. (For more detail on the project, see pages 17-18.) The project envisions working with a government partner, as well as with the Omidyar Network (which is promoting essential work in this area), industry, citizen groups, behavioral and political scientists, and others to build and pilot a system that is universal, decentralized, useful, usable, secure, privacy-enhancing, and affordable.

Participants were drawn to the transformative impact that such a project could have – not only in developing countries, but also in developed countries – as well as its ability to reduce our reliance on personally identifiable information, which is extremely vulnerable to hacks and other misuse. Among the key questions and concerns raised in the plenary discussion were whether it would be too ambitious to tackle both technological and implementation challenges (cultural, political, behavioral, etc.) inherent in a digital ID system; whether the technology required for a digital ID, such as a device capable of using biometric authentication, would be affordable at scale; how to integrate the digital ID with existing systems that are used to carry out transactions; and how to attain the buy-in of critical industries such as banking and finance. Participants noted that it would be important to understand how this project would complement, rather than replicate, efforts already underway in this field, such as those being led by the World Bank.

The second project would combine two forms of auditing to detect efforts to tamper with election tallies: a risk-limiting audit and a public audit using a technique known as end-to-end

verifiability. The former typically allows election workers to sample physical ballots and manually confirm that their contents are consistent with expectations, given the result of an election; while the latter would allow any voters, after the election, to check whether their votes were properly recorded. (For more detail on the project, see pages 13-14.) The project envisions partnering with one or more electoral jurisdictions in the United States to build and pilot a system that would apply this two-pronged auditing approach, and measure its effectiveness and impact on elections, such as voter turnout and trust. The project is rooted in the fact that no matter how hard governments and election administrators work to secure elections, they will not be impervious to tampering. And it recognizes that any attack on voting infrastructure has the potential not only to alter the results of a particular election, but also to undermine citizens' trust in the integrity of the entire electoral system, which is critical to their participation in a democracy.

Many participants highlighted the timeliness of this idea, given the increasing number of sophisticated attacks on voting systems and the declining trust of citizens in the integrity of the electoral process. It was the view of participants that, if successfully piloted, this approach would have relevance beyond the United States, as liberal democracies around the world are facing similar threats and deficits in public confidence. Among the questions raised in the plenary discussion were whether election authorities would be able to bear the costs of developing and implementing such an effort, and whether the design and testing of end-to-end verifiability technology could be achieved in the three-to-five-year window prescribed by CWP.

The third project would bring together a pair of proposals that were discussed in different working groups: information challenges in social media and privacy. (For more detail on the respective projects, see pages 20-22 and 16-17.) One proposal was to create a social media solutions laboratory, which would marshal an interdisciplinary team of software engineers, policy experts, social and behavioral scientists, journalists, and others to develop a set of online applications for social media, which would give users ways to curate and interpret the information they consume on the Internet, and mitigate challenges such as disinformation and the echo chamber effect. An application that might be among those tested in this lab could be the project proposed in the privacy working group, which operates continuously in a user's browser to create "cover traffic," making it nearly impossible for the user's real browsing data to be discerned and collected without her permission. The lab would pilot these applications with a test group of users, conducting research on their efficacy. At the same time, the lab would explore the possibility of working with one or more of the big social media platforms to develop a set of filters, which would give users greater agency in shaping their information diet.

While a few participants continued to question whether content integrity on social media should indeed be considered a cybersecurity issue (a point raised in the opening plenary session), multiple participants underscored that the challenge was central for a range of cybersecurity issues such as privacy, which will only get worse if left unaddressed. And given the current public backlash against leading social media companies due to their past failures to

address these issues, participants were of the view that they might be more willing than before to work with an outside entity on testing solutions to some of these chronic problems. That said, participants felt more work was required to articulate how the lab would actually work, including how it would be different from similar efforts to address these problems.

Finally, there were four additional proposals that – while not seen as being the right fit for full CWP projects – were, in the view of participants, areas where Columbia University and an outside partner could make a significant contribution to addressing fundamental cybersecurity challenges. It was the view of participants that these ideas, which were conceived of and developed in the context of the CWP Forum, merited further development. The first of these focused on reducing the risks associated with the supply chain for digital microelectronics (for more detail on the project, see page 14). Multiple participants pointed out that supply chain security represents a major blind spot in cybersecurity efforts, and that both government and private sector partners would benefit from applying the technology being developed at Columbia. The second project in this category proposed developing a system to facilitate real-time, automated sharing of relevant, timely, and actionable threat information to critical infrastructure, using the financial industry to develop a proof of concept (for more detail on the project, see pages 14–15). The idea was not only to enable swift and collective responses to attacks when they occur, but also to provide predictive analysis that may help prevent such attacks from taking place. Here too, the synergies developed between Columbia experts and financial industry leaders in the run up to the Forum was viewed as worth developing further. The third and fourth projects in this category were discussed in the norms and deterrence working group. One proposed bringing academic knowledge to bear on improving a set of cybersecurity norms for the tech industry, with an eye towards informing the eventual development of such norms for state actors; while the other proposed designing and testing a standard methodology to assess the effectiveness of cyber deterrence strategies pursued by the U.S. government (for more detail on these projects, see pages 26-27). Given the potential impact of these proposals and the positive feedback they received from Forum participants, CWP staff will try to facilitate the further development of these project proposals.

IV. Next Steps: Project Development, Assessment, and Implementation

For each of the three projects identified above, CWP will work with the project leads to develop a formal project proposal, which will include a description of the project, its objectives, the individuals and institutions that would be involved, a general sense of how long it would take to implement, and a rough estimate of the amount of the money it would cost. The proposals should also address outstanding questions, recommendations, and critical feedback on each project that were surfaced during the Forum. In December 2018, these three project proposals will be presented to the CWP Advisory Committee, which will advise on whether they merit further development as potential CWP projects.

Projects that are determined to merit further development will receive an initial tranche of funding to undergo a rigorous project design phase of approximately three months, during which the project leads will work with CWP staff to define the major deliverables, a precise timeline for implementation, a funding plan, a set of performance indicators for monitoring and evaluation, and the principal implementing partners. All of this information will be synthesized in a project design plan. CWP staff will then prepare an evaluation of this report, which identifies potential impact and strengths and weaknesses, and recommends whether the project should be funded. This evaluation, the project design plan, and earlier feedback from the Advisory Committee will be key factors in deciding whether these projects are implemented by CWP.

V. Acknowledgements

A number of people helped to organize and shape the CWP Forum on Cybersecurity, to whom we are profoundly grateful.

First, we thank the many individuals who gave generously of their time in advance of the Forum, informing our understanding of the challenge, developing ideas for specific projects, and suggesting potential partners for implementation. They include (unless otherwise noted, the academic institution with which individuals are affiliated is Columbia University): Alexander Abdo, Senior Staff Attorney, Knight First Amendment Institute; Jonathan Albright, Director, Digital Forensics Initiative, Tow Center for Digital Journalism; Steven M. Bellovin, Professor, Computer Science Department; Sarah Cleveland, Louis Henkin Professor of Human and Constitutional Rights and Faculty Director of the Human Rights Institute at Columbia Law School; Jared Cohen, CEO, Jigsaw; Dana DeBeauvoir, Travis County Clerk (Texas); Antonio DeSimone, Johns Hopkins University Applied Physics Lab; Doug Dicconson, Managing Partner, Theorem Studios; Renée DiResta, Data for Democracy; Jen Easterly, Managing Director of Morgan Stanley and Global Head of the Cybersecurity Fusion Center; Katy Glen Bass, Research Director, Knight First Amendment Institute; Orin Herskowitz, Executive Director, Columbia Tech Ventures; Suman Jana, Assistant Professor in Computer Science Department; Sam Jones, Palantir; Justin Kossylin, Jigsaw; Neal Kelley, Registrar of Voters for Orange County, California; C.V. Madhukar, Omidyar Network; Katherine Maher, Wikimedia Foundation; Susan McGregor, Tow Center for Digital Journalism; Ellen Meier, Director of Center for Technology & School Change and Professor of Practice, Teachers College; Fran Moore, Financial Systemic Analysis and Resilience Center; Shailagh Murray, Executive Vice President for Public Affairs; Kenneth Prewitt, Carnegie Professor of Public Affairs; Satish Rao, Associate Director of Physical Sciences Licensing, Columbia Tech Ventures; Anya Schiffrin, Lecturer in Discipline of International and Public Affairs; Bruce Schneier, Chief Technology Officer at IBM Resilient and Lecturer in Public Policy at the Harvard Kennedy School; Anthony E. Shorris, Weinberg Visiting Professor, Princeton University; Jake Sullivan, Yale Law School; Michael Ting, Professor of Political Science and International and Public Affairs; Chris Wiggins, Associate Professor in the Department of

Applied Physics and Applied Mathematics and Chief Data Scientist at the *New York Times*; Jeannette Wing, Avanessians Director of the Data Science Institute and Professor of Computer Science; Moti Yung, Google. We would like to recognize in particular two members of the Columbia faculty – Professor Salvatore Stolfo and Associate Professor Simha Sethumadhaven – both of whom developed projects that could be implemented by using intellectual property they had previously developed.

Second, we are deeply indebted to the moderators of the Forum's working groups, who went above and beyond to facilitate discussions of proposed projects and improve our work in all respects: the Hon. Jeh Johnson (critical infrastructure); Dana Hyde (privacy); Jameel Jaffer (information challenges in social media); the Hon. Lisa Monaco (standards, benchmarks, and best practices); and the Hon. John Brennan (norms and deterrence).

Third, we thank the experts who enriched the discussions of specific projects in the Forum's working groups: Alexander Abdo, Steven M. Bellovin, Sarah Cleveland, Antonio DeSimone, Doug Dicconson, Jen Easterly, Malo Hutson, C.V. Madhukar, Susan McGregor, Ellen Meier, Shailagh Murray, Hillary Schrenell, Michael Ting, Abiah Weaver, Chris Wiggins, and Moti Yung.

Fourth, we extend our appreciation to the Columbia graduate students who supported the working groups at the Forum: Veronica Nnenna Akaezuwa, Preetam Dutta, Courtney Murray, Michelle Ritter, Elizabeth Anne Watkins; to Flovijana Morina, Chandra Osborne, and Mariela Sierra from the Columbia University Programs and Events team, for their help setting up the event; to KatyAnna Johnson and Christina Shelby from the Office of the President, for their tireless assistance with invitations; and in particular to Cassie Ziegler.

Finally, our greatest thanks go to the Forum participants, whose names and biographies are listed in the annex that follows.

VI. Annex: Biographies of Forum Participants

Lee C. Bollinger
President, Columbia University

Lee C. Bollinger became Columbia University's nineteenth president in 2002. Under his leadership, Columbia stands again at the very top rank of great research universities, distinguished by comprehensive academic excellence, historic institutional development, an innovative and sustainable approach to global engagement, and unprecedented levels of alumni involvement and financial stability. Bollinger is Columbia's first Seth Low Professor of the University, a member of the Columbia Law School faculty, and one of the country's foremost First Amendment scholars. As president of the University of Michigan, Bollinger led the school's historic litigation in *Grutter v. Bollinger* and *Gratz v. Bollinger*. These Supreme Court decisions that upheld and clarified the importance of diversity as a compelling justification for affirmative action in higher education were reaffirmed in the Court's 2016 ruling in *Fisher v. University of Texas*. As Columbia's president, Bollinger conceived and led the University's most ambitious expansion in over a century with the creation of the Manhattanville campus in West Harlem. An historic community benefits agreement emerging from the city and state review process for the new campus provides Columbia's local neighborhoods with decades of investment in the community's health, education, and economic growth.

Magdi Amin
Investment Partner for Digital Identity, Omidyar Network

Magdi Amin is an Investment Partner with Omidyar Network (ON) in Washington, DC, where he focuses on global strategy and investments in Digital Identity, with a particular focus on the US and Africa. ON's vision is for a world in which everyone has access to "good identity" that empowers and protects individuals' rights to privacy and security in the digital world. Before joining Omidyar, Amin served for nearly two decades with the World Bank Group (WBG). Most recently, he managed corporate strategy at the International Finance Corporation (IFC), the private sector arm of the WBG. Previous roles included Head of Country Engagement, Principal Economist for the EVP and CEO, Manager of Investment Climate Advisory in the Middle East/North Africa (MENA) region, and Principal Economist and Strategist for the East Asia/Pacific region. In a personal capacity, he writes on Nubian cultural heritage. Amin received an A.B. from Princeton's Woodrow Wilson School and an M.A. from the Johns Hopkins School of Advanced International Studies (SAIS).

Emily Bell
Professor of Professional Practice and Director of the Tow Center for Digital Journalism, Columbia University

Emily Bell is founding director of the Tow Center for Digital Journalism at Columbia's Graduate School of Journalism and a leading thinker, commentator, and strategist on digital journalism. Established in 2010, the Tow Center has rapidly built an international reputation for research into the intersection of technology and journalism. The majority of Bell's professional career was spent at Guardian News and Media in London working as an award-winning writer and editor, both in print and online. As editor-in-chief across Guardian websites and director of digital content for Guardian News and Media, Bell led the web team in pioneering live blogging, podcasting, multimedia formats, data and social media, making *The Guardian* an internationally awarded beacon of digital transformation.

Josh Benaloh
Senior Cryptographer at Microsoft Research; Director of the International Association for Cryptologic Research

Josh Benaloh is Senior Cryptographer at Microsoft Research and an affiliate faculty member in the School of Computer Science and Engineering at the University of Washington. He earned his S.B. in mathematics from the Massachusetts Institute of Technology and his M.S., M.Phil., and Ph.D. degrees from Yale University, where his 1987 dissertation *Verifiable Secret-Ballot Elections* introduced the use of homomorphic encryption as a paradigm to enable election tallies to be verified by individual voters and observers without having to trust election equipment, vendors, or personnel. Benaloh serves on the Coordinating Committee of the Election Verification Network, spent 17 years on the Board of Directors of the International Association for Cryptologic Research, and is an author of the September 2018 report, *Securing the Vote: Protecting American Democracy*, by the National Academies of Science, Engineering, and Medicine.

Hon. John O. Brennan
Former Director of the Central Intelligence Agency

John O. Brennan served as Director of CIA (2013-2017) and Assistant to the President for Homeland Security and Counterterrorism (2009-2013). During his first stint at CIA (1980-2005), he was the daily intelligence briefer to President Clinton, Chief of Station in the Middle East, Chief of Staff to D/CIA George Tenet, Deputy Executive Director, and the first Director of the National Counterterrorism Center. Brennan has a B.A. from Fordham University and an M.A. from the University of Texas at Austin. He also studied at the American University in Cairo (1975-76).

Geoffrey Buswick
Lead Analytical Manager, U.S. Public Finance Infrastructure Group at S&P Global Ratings

Geoffrey Buswick is a Managing Director and Sector Leader in the U.S. Public Finance - Governments team at S&P Global Ratings in Boston. Buswick focuses on developing research and educating the market about S&P Global's approach to rating municipal entities and its views on emerging risks in the public finance sector. He speaks regularly on the topics of cybersecurity, sustainability, direct purchase bank loans, distressed credits, and specific aspects of the local government and school district criteria. Buswick also serves as a committee chair for local government and water & wastewater ratings. Much of his work aims at improving transparency and external understanding of S&P Global's public finance rating process. Prior to joining S&P Global, Buswick served as the Chief Financial Officer, Treasurer & Collector for the City of Gloucester, Massachussetts. In addition, he spent three years as the Administrative Officer for the City of North Adams, Massachussetts. Buswick served on the board of governors of the National Federation of Municipal Analysts (NFMA) from 2010-2013, and was the co-chair for the 2012 and 2013 NFMA Annual Conferences. In 2017, he received the "Award for Excellence" from the NFMA. Buswick also served on his town's Finance Committee from 2008-2015. Buswick holds both a B.A. in Political Science and a Masters of Public Administration from the University of Massachusetts at Amherst.

Thomas Donahue
Former Senior Director for Cyber Operations at the National Security Council

Thomas Donahue – who currently works at the environmental and energy consulting firm SC&A, Inc. and serves as a cyber advisor to The Cipher Brief – retired from CIA after 32 years of service. He served as the Chief Editor of the President's Daily Brief and other CIA daily production during the second term of the Clinton administration, and he spent the last 18 years of his career focused on cyber threats as a manager and senior analyst in what is now known as the Center for Cyber Intelligence. He served four years at the White House during the Bush and Obama administrations, most recently as the senior director for cyber operations for the National Security Council staff. During his last two years, he was the research director at the DNI's Cyber Threat Intelligence Integration Center. He has a Ph.D. in electrical engineering from the Massachusetts Institute of Technology.

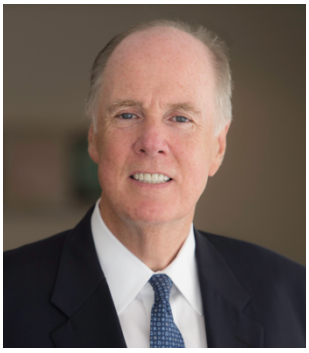

Hon. Thomas E. Donilon
Former National Security Adviser to the President; Chair of the Presidential Commission to Enhance National Cybersecurity; Chairman of the BlackRock Investment Institute

Thomas E. Donilon is Chairman of the BlackRock Investment Institute and Senior of Counsel at the international law firm of O'Melveny & Myers. He served as National Security Advisor to President Barack Obama. In that capacity, Donilon oversaw the U.S. National Security Council staff, chaired the cabinet level National Security Principals Committee, provided the President's daily national security briefing, and was responsible for the coordination and integration of the administration's foreign policy, intelligence, and military efforts. Donilon also oversaw the White House's international economics, cybersecurity, and international energy efforts. Donilon served as the President's personal emissary to a number of world leaders.

Scott Flom
IT Director, Travis County Clerk's Office, Texas

Scott Flom has spent over 20 years in Executive Management, including Operations, Product Management, and Information Technology. This includes work in the Elections Systems, Equipment Manufacturing, and IT industries. As the VP of Operations for an Elections Systems company, he had overall responsibility for the implementation of Electronic Voting Systems in two of the ten most populous counties in the United States. After work in other industries, Flom has returned to his first love, IT team management.

Ben Fried
Chief Information Officer, Google

Ben Fried is Google's Chief Information Officer, overseeing the technologies that make Googlers go. He's the New York office Tech Site Lead, responsible for 3,000 Googlers in the company's East Coast headquarters. Fried has a diverse background in systems engineering and software development. He led development of mission scheduling software for NASA at a Bay Area startup, and spent over a decade at Morgan Stanley, where he rose to the level of Managing Director, and led teams responsible for software development, internet infrastructure, and business intelligence. Fried is a graduate of Columbia University.

Nathaniel Gleicher
Head of Cybersecurity Policy at Facebook

Nathaniel Gleicher is a computer scientist and a lawyer, and works at the intersection of technology, policy, and law. He has taught computer programming, built and secured computer networks, prosecuted cybercrime at the U.S. Department of Justice, and served as Director for Cybersecurity Policy at the National Security Council (NSC) in the White House. At the NSC, he developed U.S. government policy on key technology and cybersecurity challenges, including encryption, cyber deterrence, internet governance, and network security. Since leaving government, Gleicher served as head of cybersecurity strategy at Illumio, and is currently the Head of Cybersecurity Policy at Facebook.

Donna Gregg
Head of Asymmetric Operations at the Johns Hopkins University Applied Physics Laboratory

Donna Gregg is the Sector Head for the Asymmetric Operations Sector at the Johns Hopkins University Applied Physics Laboratory (APL), where she is responsible for more than 1,300 staff supporting efforts to combat the asymmetric threats of terrorism, weapons of mass destruction, and offensive cyber. She began her career at APL in 1984 and has served in numerous leadership roles including Mission Area Executive for Cyber Operations, Managing Executive of the Applied Information Sciences Department, and Strategic Focus Area Lead for Information Assurance in the Info Centric Operations Business Area, as well as in several branch and group supervisor-level positions within the laboratory. She currently serves on the Board of Advisors for the National Institute for Hometown Security and on the Threat Reduction Advisory Committee sponsored by the Under Secretary of Defense for Acquisition, Technology & Logistics. She has an M.S. in mathematics from Johns Hopkins University and a B.S. in mathematics from the University of Maryland.

Avril Haines
Senior Research Scholar, Columbia University; Deputy Director, Columbia World Projects

Avril D. Haines is currently a Senior Research Scholar at Columbia University and a Lecturer in Law at Columbia Law School. She served as Deputy National Security Advisor to President Obama, was the Deputy Director of the Central Intelligence Agency, and served as the Legal Adviser to the National Security Council. Before joining the NSC, she led the Treaty Office at the Department of State, was the Deputy Chief Counsel for the United States Senate Committee on Foreign Relations, worked for The Hague Conference on Private International Law, and served as a law clerk for Judge Danny Boggs on the U.S. Court of Appeals for the Sixth Circuit. Haines received a bachelor's degree in Physics from the University of Chicago, a law degree from Georgetown University Law Center, and founded and ran a bookstore café for five years while engaged in community service in Baltimore.

Jason Healey
Senior Research Scholar in the School of International and Public Affairs, Columbia University

Jason Healey is Senior Research Scholar at Columbia University's School for International and Public Affairs, specializing in cyber conflict and risk. He started his career as a US Air Force intelligence officer, before moving to cyber response and policy jobs at the White House and Goldman Sachs. He was founding director for cyber issues at the Atlantic Council where he remains a Senior Fellow and is the editor of the first history of conflict in cyberspace, *A Fierce Domain: Cyber Conflict, 1986 to 2012*. He is on the DEF CON review board and served on the Defense Science Board task force on cyber deterrence.

Dana J. Hyde
Venture Partner, JVP; Former CEO of the Millennium Challenge Corporation

Dana Hyde is a senior executive with over 25 years of experience in law, public policy, and international development. She served eight years in the Obama Administration, as Chief Executive Officer of the Millennium Challenge Corporation (MCC), Associate Director at the White House Office of Management and Budget (OMB), and Senior Advisor to the Deputy Secretary of State. Hyde also served as Counsel to the 9/11 Commission and as Special Assistant to the Deputy Attorney General in the Clinton Administration. Earlier in her career Hyde practiced law at WilmerHale in London and in Washington, DC at Zuckerman Spaeder. She is currently a Partner at the venture capital firm JVP.

Jameel Jaffer
Executive Director of the Knight First Amendment Institute at Columbia University

Jameel Jaffer is inaugural director of the Knight First Amendment Institute at Columbia University. Previously, he was deputy legal director of the American Civil Liberties Union, where he oversaw that organization's work relating to free speech, privacy, national security, and international human rights. He has argued cases at all levels of the federal court system, including in the Supreme Court, and has testified many times before Congress and administrative agencies. His most recent book, *The Drone Memos*, was published by the New Press in the fall of 2016.

Merit Janow
Dean of the School of International and Public Affairs, Columbia University

Merit E. Janow is an internationally recognized expert in international trade and investment, with extensive experience in academia, government, international organizations, and business. She is Dean of the Faculty and a Professor of Practice in International Economic Law & International Affairs at Columbia University's School of International and Public Affairs (SIPA) and affiliated faculty at Columbia Law School. As Dean, she has launched a major initiative around technology and policy, and initiated new research and programs around the digital economy, cyber security, and internet governance. Janow has had three periods in government service: in December 2003, Janow was elected for a four-year term as one of the seven Members of the World Trade Organization's (WTO) Appellate Body – the first female to serve as a judge on the Geneva based appellate body, which hears government to government disputes on economic issues. From 1997 to 2000, she served as the Executive Director of the first international antitrust advisory committee of the U.S. Department of Justice. Prior to joining Columbia's faculty, she was Deputy Assistant U.S. Trade Representative for Japan and China (1989-93), responsible for all bilateral trade negotiations between the United States and both Japan and China. She negotiated more than a dozen trade agreements. She has written two books and numerous articles. Janow is on the Board of Directors of several companies in financial services and technology and several not-for-profit organizations such as MasterCard, the American Funds, and the National Committee on United States-China Relations. She was raised in Tokyo, Japan and speaks Japanese.

Sam Jeffers
Co-founder of Who Targets Me

Jeffers is the co-founder of Who Targets Me (https://whotargets.me), software to help voters understand how political campaigns are chasing their vote using social media advertising. To date, over 10,000 people have installed Who Targets Me in more than 50 countries. The project has tracked political advertising in 10 election campaigns in the last 18 months and is currently monitoring its use in the Brazilian and US midterm election campaigns.

Eric Johnson
Professor of Business at Columbia Business School and Director of the Center for the Decision Sciences

Eric Johnson is the Norman Eig Chair of Business and Director of the Center for Decision Sciences at Columbia Business School. His research examines the interface between behavioral decision research, economics and the decisions made by consumers, managers, and their implications for public policy, markets, and marketing. He was awarded the Distinguished Scientific Contribution Award from the Society for Consumer Psychology, and named a Fellow by the Association for Consumer Research, was awarded an honorary doctorate in Economics from the University of St. Gallen, and is a Fellow of the Association for Psychological Science. According to the Institute for Scientific Information, he is one of the most highly cited scholars in Business and Economics. He served as a senior visiting scholar at the Consumer Financial Protection Bureau.

Hon. Jeh Johnson
Former Secretary of the U.S. Department of Homeland Security

Jeh Johnson is a partner with the law firm Paul, Weiss, Rifkind, Wharton & Garrison LLP and the former U.S. Secretary of Homeland Security (2013-2017). Prior to that, Johnson was General Counsel of the Department of Defense (2009-2012), General Counsel of the Department of the Air Force (1998-2001), and an Assistant United States Attorney for the Southern District of New York (1989-1991). Johnson is a Fellow in the American College of Trial Lawyers and a member of the Council on Foreign Relations. He is a graduate of Morehouse College (1979) and Columbia Law School (1982), and the recipient of nine honorary degrees. Johnson frequently lectures at Harvard, Yale and other law schools, and is a non-resident Senior Fellow at the Harvard Kennedy School.

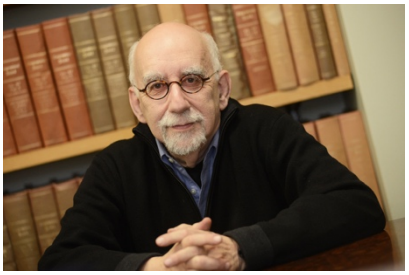

Ira Katznelson
Ruggles Professor of Political Science and History

Ira Katznelson is Ruggles Professor of Political Science and History at Columbia University. His 2013 book *Fear Itself: The New Deal and the Origins of Our Time* has been awarded the Bancroft Prize in History and the Woodrow Wilson Foundation Award in Political Science. Other books include the just-published *Southern Nation: Congress and White Supremacy After Reconstruction* (co-authored with David Bateman and John Lapinski). Katznelson is a former president both of the American Political Science Association and the Social Science Research Council. He earned his B.A. at Columbia College and his Ph.D. in History at the University of Cambridge, where he served in 2017-18 as Pitt Professor of American History and Institutions.

Nicholas Lemann
Director, Columbia World Projects; Director, Columbia Global Reports; Joseph Pulitzer II and Edith Pulitzer Moore Professor of Journalism; Dean Emeritus of the Faculty of Journalism

Nicholas Lemann directs Columbia World Projects, a new initiative to connect academic work with entities beyond the academy that possess the power and influence to transform research into concrete consequences benefiting humanity. He also directs Columbia Global Reports, a book publishing venture that presents reporting around the globe on a wide range of political, financial, scientific, and cultural topics. Lemann is Dean Emeritus and Pulitzer Moore Professor of Journalism at Columbia. During his deanship, the Journalism School completed its first capital fundraising campaign, started its first new professional degree program since the 1930s, and launched significant initiatives in investigative reporting, digital journalism, and executive leadership for news organizations. Board memberships include Columbia's Knight First Amendment Institute and the Russell Sage Foundation. Lemann is a member of the New York Institute for the Humanities and the American Academy of Arts and Sciences, and a staff writer for *The New Yorker*.

Charles Luftig
Policy and Standards Lead, Bridgewater Associates

Charles Luftig is the Head of Security Policy for Bridgewater Associates. Prior to working at Bridgewater, he served as the Deputy General Counsel at the Office of Management and Budget at the White House. He also held several legal and policy roles at the National Security Council as Senior Adviser to the Deputy National Security Adviser (2015-2016), Deputy Legal Adviser to the National Security Council (2014-2015), and Director for Counterterrorism (2012-2013). Luftig also worked in the National Security Division at the Department of Justice from 2010-2012, where he received the 2013 Attorney General Award for Excellence (the highest honor awarded by the Department) and the 2012 Assistant Attorney General Award for Excellence (the highest honor awarded by the Division). Earlier in his career, he worked as a litigator in private practice and clerked for the Honorable Colleen Kollar-Kotelly in the United States District Court for the District of Columbia. Luftig is a graduate of the University of Michigan and the University of Virginia School of Law.

Michael Lynton
Chairman of Snap, Inc.

Michael Lynton served as the CEO of Sony Entertainment from April 2012 until February 2017, overseeing Sony's global entertainment businesses, including Sony Music Entertainment, Sony/ATV Music Publishing and Sony Pictures Entertainment. Beginning in January 2004, Lynton also served as Chairman and CEO of Sony Pictures Entertainment. Prior to joining Sony Pictures, Lynton worked for Time Warner and served as CEO of AOL Europe, President of AOL International, and President of Time Warner International. He earlier served as Chairman and CEO of Pearson PLC's Penguin Group where he oversaw the acquisition of Putnam, Inc. and extended the Penguin brand to music and the Internet.

Lynton currently serves on the Board of Snap, Inc., Pearson, PLC, Schrodinger, LLC, IEX, and Ares Management, L.P. He is also a member of the Council on Foreign Relations and the Harvard Board of Overseers, and serves on the boards of the Los Angeles County Museum of Art, the Tate, and the Rand Corporation. Michael holds a B.A. in History and Literature from Harvard College, where he also received his M.B.A.

John J. MacWilliams
Fellow at the Center on Global Energy Policy, Columbia University; Former Associate Deputy Secretary of the U.S. Department of Energy

John J. MacWilliams is a Fellow at the Center on Global Energy Policy at Columbia University. Prior to joining CGEP, MacWilliams served as Associate Deputy Secretary of the U.S. Department of Energy after being appointed in August 2015. He also served as DOE's Chief Risk Officer and, before that, as a Senior Advisor to the Secretary. Prior to DOE, MacWilliams was a partner of Tremont Energy Partners, LLC, a private investment firm based in Cambridge, Massachusetts. Prior to Tremont, he was Vice Chairman, Investment Banking, at JP Morgan Chase, a Partner of JP Morgan Partners, and a founding partner of The Beacon Group, LLC, a private investment firm located in New York and acquired by JPMorgan Chase in 2000. He was also Partner and Co-Head of the Beacon Group Energy Investment Funds. Prior to the formation of The Beacon Group, MacWilliams was with Goldman Sachs & Co. and an attorney at Davis Polk & Wardwell. MacWilliams holds a BA from Stanford, an MS from MIT, and a JD from Harvard Law School.

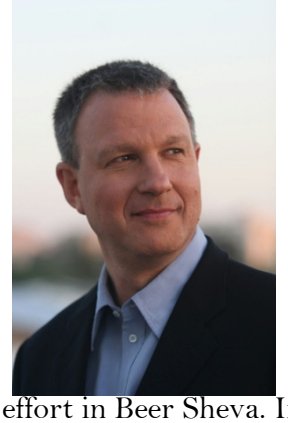

Erel Margalit
Founder and Chairman of JVP

Erel N. Margalit, Ph.D., founded JVP over 24 years ago, and is one of the chief architects of the startup nation, bringing innovation and entrepreneurial leadership to the country's most pressing political, economic, and social challenges. Chosen for the Forbes Midas List as the first venture capitalist with the golden touch, Margalit has led numerous global exits as managing partner of JVP, including implementing the first state sponsored incubator into the famed JVP family of funds, creating the Media Quarter in Jerusalem, and leading the Cyber Security effort in Beer Sheva. In addition to helping orchestrate the $4.8 billion sale of Chromatis to Lucent Technologies in 2000, Margalit successfully led investments in companies such as QlikTech (NASDAQ:QLIK), Cogent Communications (NASDAQ: CCOI), Netro (initial public offering; JVP subsequently sold its shares at a $5.5 billion company valuation), Cyoptics (acquired by Avago for $434 million), Precise (initial public offering and subsequently acquired by Veritas), Scorpio (acquired by US Robotics), Fundtech, ViryaNet, Jacada, and Allot, among others, in addition to leading the investment and serving as the Chairman of CyberArk Software (NASDAQ: CYBR) for numerous years. During his tenure as Member of the Israeli Knesset from 2013-2017, Margalit served as a Member of the Security and Foreign Affairs Committee and Finance Committee. Margalit led the Knesset's Cybersecurity Taskforce and the Taskforce developing the North and South of Israel, implementing various economic development initiatives around the country. Margalit also spearheads the Innovation Initiative for the Mediterranean Basin countries, engaging 14 Arab countries on projects around innovation and cyber security cooperation in the Euro-Med region.

Mike Masnick
Founder and CEO of Floor64; Editor of the Techdirt Blog

Mike Masnick is the founder & editor of the popular Techdirt blog as well as the founder of the Silicon Valley think tank, the Copia Institute. In both roles, he explores the intersection of technology, innovation, policy, law, civil liberties, and economics. His writings have been cited by Congress and the EU Parliament. According to a Harvard Berkman Center study, his coverage of the SOPA copyright bill made Techdirt the most linked-to media source throughout the course of that debate. Masnick and Techdirt have also been key players in the ongoing battles over net neutrality and encryption. Masnick is also known for coining the term "The Streisand Effect," to describe how attempting to stifle speech online can serve to draw even more attention.

Hon. Denis McDonough
Visiting Senior Fellow in Carnegie's Technology and International Affairs Program; Former Chief of Staff to the President of the United States

Denis McDonough served as White House Chief of Staff to President Barack Obama from February 1, 2013 to January 20, 2017. In this role, he managed a four-thousand-member White House staff, as well as Cabinet Secretaries and agency leaders. He provided strategic advice to the President on the most significant domestic policy, national security, and management issues facing the federal government; enforced plans and accountability for performance goals; and planned and coordinated efforts to recruit and retain key talent in the federal government. Prior to the White House, McDonough served in senior leadership and policy-making positions in the U.S. House of Representatives and U.S. Senate. McDonough is currently a Senior Principal at the Markle Foundation and chairs its Rework America Task Force, a national initiative to transform the labor market so that all Americans can thrive in the digital economy. He also serves as an executive fellow at the University of Notre Dame's Keough School of Global Affairs. McDonough is a graduate of St. John's University (MN) and Georgetown University School of Foreign Service.

James Miller
President and CEO of Adaptive Strategies LLC; Former Under Secretary of Defense for Policy

Dr. James N. Miller is President and CEO of Adaptive Strategies LLC, which advises clients on technology issues and building organizational capacity. He is a senior fellow at Johns Hopkins University's Applied Physics Lab and at Harvard University's Belfer Center for Science & International Affairs. Miller serves on the Board of Advisors for Endgame, Inc. and for the Center for a New American Security. He is a member of the Council on Foreign Relations, the International Institute for Strategic Studies, and the Defense Science Board, where he co-chaired a study on cyber deterrence. He served as Under Secretary of Defense for Policy in the Obama Administration, where he led the development of a national defense strategy for cyberspace.

Hon. Lisa Monaco
Former Assistant to the President for Homeland Security and Counterterrorism; Co-Chair of the Aspen Institute Cyber Group

Lisa Monaco has spent more than two decades in public service and in senior management and advisory positions in law enforcement and national security. As the President's Homeland Security and Counterterrorism Advisor from 2013-2017, she coordinated the federal government's crisis management and response to a wide array of risks and challenges, from cyberattacks and pandemics to terrorist threats. Monaco also spent 15 years at the Department of Justice, serving as a career federal prosecutor, as well as in senior management positions in the Justice Department and the FBI, where she served as Chief of Staff to then-FBI Director Robert S. Mueller, III and helped him lead the FBI's post-9/11 transformation. In 2011, she was nominated and confirmed to serve as Assistant Attorney General for National Security, the first woman to serve in that position. In this role, she oversaw all federal terrorism and national security prosecutions. She made investigating and prosecuting national security cyber threats a top priority and created a nationwide network of national security cyber prosecutors. Monaco is a graduate of Harvard University and the University of Chicago Law School.

Kate O'Sullivan
General Manager, Microsoft

Kate O'Sullivan is the general manager of Digital Diplomacy within Microsoft's Corporate, External and Legal Affairs Department. In this role, O'Sullivan leads Microsoft's global cyber security strategy and related public policy campaigns such as its Defending Democracy and Digital Peace initiatives. Prior to working at Microsoft, O'Sullivan was managing director at the global communications and public affairs consultancy, Burson-Marsteller, heading up their Corporate Practice in San Francisco. Earlier in her career, she was on the co-founding team and was vice president of marketing and communications at SEVEN Networks. Before SEVEN, O'Sullivan was vice president and managing director at Ogilvy Public Relations Worldwide, and managed Ogilvy's Silicon Valley technology practice. O'Sullivan currently serves on the board of 826 National, IESE Business School, Bay Area Council and sf.citi. She graduated from University of Colorado in Boulder with a degree in International Relations.

Sean Roche
Associate Deputy Director of CIA for Digital Innovation

Sean P. Roche is a senior government executive with more than 37 years of federal service. In March 2015, he was named Associate Deputy Director of CIA for Digital Innovation, serving as the second-in-command of CIA's first new directorate in more than 50 years. The Directorate of Digital Innovation (DDI) is responsible for cyber intelligence, open source collection, secure global communications, worldwide mission information systems, data curation, and data science. The DDI also accelerates the integration of advanced digital capability across all of CIA's mission areas. Additionally, the CIA's Chief Information Officer (CIO) and Chief Data Officer (CDO) serve within the DDI. Immediately prior to this position, Roche served as the Associate Deputy Director of CIA for Science and Technology. Over the course of his career, Roche has held various senior leadership positions across a wide range of missions, disciplines, and tradecraft, ranging from research and development to clandestine operations. Within the Directorate of Science and Technology, he served in senior leadership roles in the Offices of Development and Engineering, Technical Collections, Global Access, Integrated Missions, and Mission Resources. He led teams that developed, delivered, and deployed satellite and airborne reconnaissance systems, next generation collection platforms, clandestine technical operations, and advanced targeting tradecraft. In addition to his assignments at CIA, Roche has also served in positions across the Intelligence Community and the Department of Defense. He is a recipient of the Distinguished and Meritorious Presidential Rank Awards, the CIA Director's Award, the Directorate of Operations Donovan Award, and other Intelligence Community and CIA Meritorious Unit Citations. He became a member of the Senior Intelligence Service in June 2001.

Henning Schulzrinne
Professor in the Computer Science Department, Columbia University; Former Chief Technology Officer for the United States Federal Communications Commission

Prof. Henning Schulzrinne, Levi Professor of Computer Science at Columbia University, received his Ph.D. from the University of Massachusetts in Amherst, Massachusetts. He was a member of technical staff at AT&T Bell Laboratories, Murray Hill and an associate department head at GMD-Fokus (Berlin), before joining the Computer Science and Electrical Engineering departments at Columbia University. He served as chair of Computer Science Department from 2004 to 2009, and as Engineering Fellow, Technical Advisor, and Chief Technology Officer of the Federal Communications Commission (FCC) from 2010 until 2017. Protocol standards co-developed by Schulzrinne, including RTP, RTSP and SIP, are now used by almost all Internet telephony and multimedia applications. He is a fellow of the ACM and IEEE.

Simha Sethumadhaven
Associate Professor of Computer Science, Columbia University

Simha Sethumadhavan is an Associate Professor of Computer Science at Columbia University. Sethumadhavan's research at Columbia is focused on finding practical solutions to problems in the areas of cybersecurity and computer architecture. He is a recipient of an Alfred P. Sloan Research Fellowship, the NSF CAREER award and an IBM co-operative research award. His work has received nine best paper awards for his work on computer security and computer architecture, and his team has successfully taped out three novel computing chips (e.g., an analog-digital computing chip) on shoestring budgets. His team's work on identifying security vulnerabilities resulted in fixes to major products such as mobile phone processors and web browsers used by millions of users, and his work on hardware security is actively considered by standards organizations. He has served on the Federal Communications Commission Downloadable Security Technical Advisory Committee. He is the Founder & CEO of Chip Scan Inc., a hardware security company focused on finding and mitigating hardware backdoors. Sethumadhavan obtained his Ph.D. from the University of Texas at Austin in 2007.

Brad Smith
President and Chief Legal Officer of Microsoft

Brad Smith serves as the President of Microsoft Corporation. He leads a team of 1,500 professionals working in 56 countries, where they are responsible for the company's legal work, intellectual property portfolio, corporate philanthropy, public policy, corporate governance, social responsibility issues and compliance matters. Smith plays a key role in representing the company externally and in leading its work on a number of critical issues including privacy, security, accessibility, environmental sustainability, human rights and digital inclusion. In 2013, he was named by the *National Law Journal* as one of the 100 most influential lawyers in the United States. In 2014, the *New York Times* called Smith "a de facto ambassador for the technology industry at large."

Alissa Starzak
Director of Public Policy, Cloudflare

Alissa Starzak is the Head of Policy at Cloudflare, a web security and optimization company. Prior to joining Cloudflare, Starzak worked for the U.S government in a variety of national security positions, including serving as General Counsel of the U.S. Army, Deputy General Counsel (Legislation) at the Department of Defense, counsel to the Senate Select Committee on Intelligence, and assistant general counsel at the Central Intelligence Agency. She also worked in private practice in Washington, D.C., and clerked for The Honorable E. Grady Jolly, U.S. Court of Appeals for the Fifth Circuit. She graduated from Amherst College and the University of Chicago Law School, where she served as an editor of the University of Chicago Law Review.

Nik Steinberg
Forum Director, Columbia World Projects

Nik Steinberg is the Forum Director at Columbia World Projects. He previously served as the Counselor and Chief Speechwriter for Amb. Samantha Power, U.S. Ambassador to the United Nations. Prior to that, Steinberg was Senior Researcher in the Americas Division of Human Rights Watch, where his work focused primarily on Mexico and Cuba. He is a graduate of Dartmouth College and the Harvard Kennedy School of Government.

Salvatore Stolfo
Professor of Computer Science, Columbia University

Salvatore Stolfo has been a Professor of Computer Science at Columbia University since 1979. He is credited as creating the area of machine learning applied to intrusion detection in cybersecurity and has created several anomaly detection algorithms and systems addressing some of the hardest problems in securing computer systems. He was elevated to IEEE Fellow for his contributions to machine learning based computer security. Stolfo has been a member of several National Academy committees evaluating Army and Navy cybersecurity activities. He has been granted over 80 patents, several of which have been engaged in major infringement litigation. Stolfo has published hundreds of papers, with many receiving best paper awards. He has consulted for major U.S. financial institutions, U.S. government agencies, defense contractors, and serves as an advisor to a venture capital firm. Two cybersecurity companies, Allure Security and Red Balloon Security, have been spun out from his Intrusion Detection Laboratory. Both companies have been actively sponsored by the Department of Defense and Department of Homeland Security.

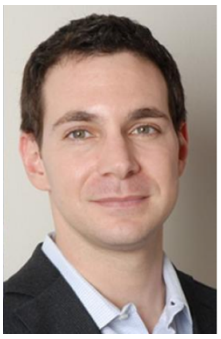

Matthew Waxman
Liviu Librescu Professor of Law, Columbia Law School

Matthew C. Waxman is the Liviu Librescu Professor of Law at Columbia Law School, where he directs the national security law program. He is also co-chair of the Cybersecurity Center at Columbia University's Data Science Institute and Adjunct Senior Fellow for Law and Foreign Policy at the Council on Foreign Relations.

Waxman previously served at the U.S. Department of State, as Principal Deputy Director and Acting Director of the Secretary of State's Policy Planning Staff. His prior government appointments included Deputy Assistant Secretary of Defense, Director for Contingency Planning and International Justice at the National Security Council, and executive assistant to National Security Advisor Condoleezza Rice. Waxman is a graduate of Yale College and Yale Law School, and he studied international relations as a Fulbright Scholar in the United Kingdom. After law school, he served as law clerk to Supreme Court Justice David H. Souter and U.S. Court of Appeals Judge Joel M. Flaum. Earlier in his career he worked as an analyst at RAND.

Jeannette M. Wing
Avanessians Director of the Data Science Institute and Professor of Computer Science, Columbia University

Jeannette M. Wing is Avanessians Director of the Data Science Institute and Professor of Computer Science at Columbia University. From 2013 to 2017, she was a Corporate Vice President of Microsoft Research. She is Consulting Professor of Computer Science at Carnegie Mellon where she twice served as the Head of the Computer Science Department and had been on the faculty since 1985. From 2007-2010 she was the Assistant Director of the Computer and Information Science and Engineering Directorate at the National Science Foundation. She received her S.B., S.M., and Ph.D. degrees in Computer Science, all from the Massachusetts Institute of Technology. Wing's general research interests are in the areas of trustworthy computing, specification and verification, concurrent and distributed systems, programming languages, and software engineering. She received Distinguished Service Awards from the CRA and ACM. She is a Fellow of the American Academy of Arts and Sciences, AAAS, ACM, and IEEE.

Major General (ret.) Amos Yadlin
Director of the Institute for National Security Studies (INSS), Tel Aviv University

Major General (ret.) Amos Yadlin has been the Director of Tel Aviv University's Institute for National Security Studies (INSS), Israel's leading strategic think tank, since November 2011. Yadlin was designated Minister of Defense of the Zionist Union Party in the March 2015 elections. Yadlin served for over 40 years in the Israel Defense Forces, nine of which as a member of the IDF General Staff. From 2006-2010, Yadlin served as the IDF's chief of Defense Intelligence. From 2004-2006, he served as the IDF attaché to the United States. In February 2002, he earned the rank of major general and was named commander of the IDF Military Colleges and the National Defense College. A former deputy commander of the Israel Air Force, Yadlin has commanded two fighter squadrons and two airbases. He has also served as Head of IAF Planning Department (1990-1993). He accumulated about 5,000 flight hours and flew more than 250 combat missions behind enemy lines. He participated in the Yom Kippur War (1973), Operation Peace for Galilee (1982) and Operation Tamuz – the destruction of the Osirak nuclear reactor in Iraq (1981). Yadlin holds a B.A. in economics and business administration from Ben-Gurion University of the Negev (1985). He also holds a Master's degree in Public Administration from the John F. Kennedy School of Government at Harvard University (1994).