Center for the Study of WMD Weapons of Mass Destruction
NATIONAL DEFENSE UNIVERSITY

wmdcenter.ndu.edu
PROCEEDINGS

How Emerging Technologies Become Emerging Threats: Workshop Report

April 2023

By Gerald L. Epstein, Diane DiEuliis, and Quinn Urich

On October 27, 2022, the Center for the Study of Weapons of Mass Destruction, a component in the Institute for National Strategic Studies at the National Defense University, hosted an analytic workshop to discuss how emerging technologies become emerging threats. At the request of the Defense Threat Reduction Agency's Operations and Integration Directorate/Operational Analysis Department, the workshop brought together several U.S. Governmental and external experts in the areas of technology watch, horizon-scanning, and technology assessment.

Executive Summary

Today, the United States struggles with identifying and prioritizing the plethora of potential vulnerabilities to national security posed by emerging technologies. Emerging technologies are often confused with emerging threats, but the two categorizations are not synonymous. Identifying which emerging technologies indeed constitute emerging threats can better prepare society to take the appropriate actions to mitigate hazards and possibly introduce governance measures that ensure better control over their development. The participants in the workshop described in this paper found that social, cultural, political, economic, and other factors are what make emerging technologies emerging threats. The confusion often has its roots in a "hype cycle" that stresses the dangers of an emerging technology while ignoring the various forms of expertise and resources (personnel and otherwise) needed to implement it. Denying bad actors important technical expertise or tacit knowledge plays an important role in denying them access to emerging technologies that may pose a threat to society.

Note that this workshop addressed only the possible hostile exploitation of emerging technologies by adversaries. New technologies can also cause harm through accidents or through unforeseen consequences, particularly ones that are indirect. Such consequences should be anticipated and mitigated as well, but they were not the focus of this workshop.

The participants found that even exploring the capabilities of emerging technologies can pose a risk to national security. Researchers beginning investigations in a certain area run the risk of signaling intent, motivating adversaries to respond with their own development programs. On the other hand, a lack of public research on a particular topic may have the same effect because adversaries may interpret that gap as evidence that a classified research program exists within the U.S. Government. Participants agreed that vulnerabilities often need to be publicized to generate

Gerald L. Epstein is a Distinguished Fellow in the Center for the Study of Weapons of Mass Destruction (CSWMD), Institute for National Strategic Studies, at the National Defense University. Diane DiEullis is Assistant Director and Distinguished Research Fellow in CSWMD. Quinn Urich is a Research Trainee in CSWMD.

support to address them—but the need to motivate action to mitigate them is in tension with the risk that these warnings may motivate adversaries to exploit them.

During the workshop, participants pointed to different values and ethics among cultures over technology development. Chinese and U.S. researchers prioritize ethical guidance on artificial intelligence differently, for example. And within the West, Europe and the United States have different priorities and approaches to some technologies, such as genetically modified crops, data mining, and artificial intelligence. Participants agreed that a common lexicon is essential to the development of international norms and that establishing the language on an issue is important to setting standards that can underlie norms or give one party or another an economic advantage. The near infinite number of vulnerabilities to potential technological threats forces decisionmakers to prioritize these threats by placing them on a spectrum that considers what adversaries might find easier or harder to implement, and these considerations include human factors.

Ultimately, the workshop found that there is no simple formula for determining what makes an emerging technology an emerging threat; factors such as their funding, their reliance on tacit knowledge, their institutional support/infrastructure, their use, and controls or governance approaches over them all contribute. Workshop participants agreed that wargames or interactive simulations could be used to explore which government structures are most effective at preventing emerging technologies from becoming emerging threats. Case studies of past emerging technologies/emerging threats offer a wealth of information for researchers and allow them to better understand potential indicators, threat dynamics, and governance models.

Background and Description

Background

On October 27, 2022, the National Defense University's Center for the Study of Weapons of Mass Destruction hosted a workshop to examine how emerging technologies become emerging threats. This workshop was conducted at the request of the Defense Threat Reduction Agency, and it brought together several U.S. Governmental and external experts in technology watch, horizon scanning, and technology assessment. It was conducted subject to the National Defense University's nonattribution rule: information discussed during the workshop could be used freely, but unless the speaker explicitly authorized otherwise, nothing said at the workshop was to be attributed directly or indirectly to the speaker in the presence of anyone not authorized to participate. Accordingly, this report summarizes the workshop's discussions, but it does not attribute any of the remarks. Unless otherwise explicitly stated, none of the material included here should be considered a consensus point among workshop participants.

Purpose

Technology watch and horizon-scanning activities are conducted to provide warning of technological developments that might create adverse consequences—whether used by adversaries to the detriment of the security of the United States or its allies and partners, or in commercial applications that might introduce new societal vulnerabilities. The warning that such forecasting and assessment activities provide would ideally allow measures to be taken to mitigate those adverse consequences, and it might even highlight the need for and possibility of governance measures that could be applied to the technology's development and/or utilization that would make it less likely to pose problems.

For the workshop, an *emerging threat* was defined as the confluence of nascent *technical capability* with *vulnerability* and *capable actors* in ways that resist *mitigation*. *Technical capability* was defined as the ability to make something happen. *Vulnerability* was used to describes the adverse consequences of that something's occurrence. Vulnerabilities can include liabilities particular to the military, to the broader national security and government apparatus, or to society at large. *Capable actors* were defined as those parties willing and able to exploit the nascent technology, with their ability to do so in turn depending on their expertise, resources, interest/intent, and commitment. *Mitigation* was defined as representing how easily such a plan could be deflected, deterred, detected, impeded, or countered; among other things, mitigation can be a function of chokepoints in the process of technological development or utilization that provide opportunities for monitoring or control.

Although the initial objective of the workshop was to address the utilization of new technologies by adversaries that would provide them with some sort of military or security advantage or the ability to inflict harm, workshop participants also realized that the commercial application of new technologies by those who do not have malicious intent could also introduce vulnerabilities that could be exploited by those who do. A good example is the abuse of social media to wage disinformation campaigns. Such consequences can also be considered emerging threats.

One example of assessment of the threat posed by emerging technologies is the 2018 report of the National Academies of Sciences, Engineering, and Medicine titled *Biodefense in an Age of Synthetic Biology*, which provides a framework with which to assess the level of concern that should be attached to various developments arising from synthetic biology.[1] This framework addresses the following factors:

- usability of the technology
- usability as a weapon
- requirements of actors
- potential for mitigation.

The "usability as a weapon" criterion refers to, but does not extensively address, what Kathleen M. Vogel and Sonia Ben Ouagrham-Gormley call "socio-technical factors" that affect the ability of malefactors to use new technologies to generate harm. In a recent paper, they state:

> *For far too long, the U.S. security community has jumped to assumptions about how easily data or materials (usually related to emerging technologies) can be translated to security threats. Usually, the standard pieces of evidence used to make those judgments are generic references to the technology (or assumed trends of the technology) without relying on rigorous, real-world empirical data and studies of the various social and technical factors involved in shaping the development and use of that data or technology.*[2]

At the workshop, Vogel elaborated that there are a variety of social, cultural, political, economic, management, and organizational factors that must be brought to bear to design, develop, and deploy a technology to work in practice. Past examples of state and nonstate actor attempts to capitalize on new technology indicate that any one of these factors, if not fulfilled, can lead to failure even if the requisite data, materials, and technical expertise are all present.

Discussion

In a broadly ranging discussion, a few themes emerged.

Emerging technology often is conflated with *emerging threat*, and indeed, differentiating the two for national security purposes was the objective of the workshop. Emerging technologies are not necessarily threatening. Social, cultural, political, economic, and other factors are what make emerging technologies emerging threats. Because regulation, legislation, development of countermeasures, and other mitigation measures often lag the threat, it is important to receive warning of actual threats that may emerge. Ten years from now, what will we wish we had done today? However, because there are so many emerging technologies and they are developing so quickly, prioritization is necessary—and hence the need to determine which of these emerging technologies will most probably pose threats.

For example, different factors go into creating a bioweapons threat, including agents, procedures, and technical expertise. An individual may have the materials for a bioweapon but lack the scientific knowledge to produce and disseminate it effectively. Accordingly, the technical capabilities necessary to develop a weapon are often the subject of debate. However, even if technical knowledge that can be exploited by bad actors is available, the process of acquiring and being able to act on that knowledge—especially when it involves tacit aspects that are not readily codifiable or transferable—and other social factors may be minimized or ignored by those assessing the threat. The "knowledge-transfer problem" can mitigate some of the risk. Assessing these emerging threats requires a multidisciplinary approach consisting of experts from scientific, technical, engineering, and mathematical fields; historians; social scientists; and regional/area experts.

Tacit Knowledge

Assessing how threatening a technological development might be requires understanding the importance of tacit knowledge in bringing the technology to fruition. Even the ability to procure parts of a technology development process commercially—eliminating the need to develop or replicate it—requires solving knowledge-transfer problems. Scientific and technological know-how do not transfer easily.

Denial of small amounts of tacit knowledge can hinder technological development. For example, in 2002, researchers at the State University of New York published a description of their synthesis of infectious poliovirus from scratch.[3] The materials, genomic information, and techniques to do this were thus made publicly available, but other researchers had difficulty reproducing the results because a certain step was highly dependent on tacit knowledge. Although this paper was published 20 years ago, the experiment remains as difficult to reproduce today as it was then, indicating that the march of time and technology does not necessarily eliminate some of the barriers to replication that a lack of tacit knowledge creates.

On the other hand, just as a lack of tacit knowledge can impede a weapons development program, the acquisition of tacit knowledge in related but benign areas can contribute to such a program's success. For example, prospects for technology cooperation with potential adversaries must be evaluated regarding not only the technology's direct application but also other purposes to which it may be applied—with the recognition that the further from the domain in which the technology has been demonstrated, the greater the knowledge-transfer problem.

The constraints and opportunities provided by tacit knowledge change over time. Some tacit knowledge barriers are overcome as technology advances and diffuses—what might have been a highly tacit-knowledge-dependent procedure at one point may later become widely taught or obviated by advances in other technologies. Previously tacit knowledge might be codified in automated processes. For example, a laboratory robot could be programmed to do manipulations that an experienced technician would otherwise be needed to perform. Even if it proved challenging to capture the technician's actions precisely, once that had been accomplished, any such robot could reproduce the procedures exactly. On the other hand, some tacit knowledge may have a sensory component that can be acquired only through practice or may be embodied in communal forms that require teams to be able to acquire and transfer.

A participant in the workshop explained that tacit knowledge requirements can prevent individuals unfamiliar with a particular technology from readily acquiring it, but they are less constraining for individuals who are already working in the field. For example, more than a dozen academic and industrial labs today are conducting research in areas such as enhancing the transmissibility of flu viruses that are much more lethal than COVID-19, modifying measles to escape immune system defenses, or developing step-by-step protocols to engineer SARS-CoV-2 variants that can evade immune system defenses—research that, when published, could make these techniques accessible to yet more trained professionals. On the other hand, if teams are involved—as they typically are—it is not the tacit knowledge of one individual but of the team collectively that matters. Some observers state that sufficiently motivated researchers working in one of these labs could synthesize a respiratory virus, infect themselves, and trigger a pandemic without having to master skills not typically used in such a laboratory, such as aerosolization and dissemination. Others, however, argue that we still do not understand enough about how viruses work in nature, their mechanisms of virulence and pathogenesis, and the way they propagate to be confident that such an approach would work as intended.

The Hype Cycle—and Looking Past It

Hype and expectations regarding new technological developments go through a familiar cycle. Initially, there is shock and amazement at the promise of a new technology—or at the dangers it may pose. (If you can hype a technology, you can also be a fearmonger.) However, this hype often ignores the expertise and other factors necessary to use the technology effectively in ways that could pose a threat. Working-level perspectives among those involved in developing the technology usually show that there is more complexity involved, and that the necessary skill set is more difficult and fraught with challenges than the public hype would imply. People at the working level may understand these complexities better than senior academic principal investigators or industry representatives.

Overlooking potential issues related to the development of each new technology allows the technology to be caught up in the next hype cycle. For example, a failed hype cycle concerned nuclear energy in the United States in the 1950s. During that time, nuclear energy was seen as destined to power everything from cars and planes to the entire U.S. electricity grid. However, political, social, and industry factors all played a role—along with technical ones—in nuclear power's not living up to its expectations. Focusing on materials and expertise is a good—and accessible—starting point, but it is not enough. A broad, multidisciplinary approach that goes beyond strictly technical aspects is necessary.

Emergence of a technology can result from a combination of fortuitous circumstances, as well as prolonged investment and inquiry—engineering ability, financial will, political will, and the right kind of people. It may also be the result of new capabilities finally making possible developments that had been anticipated—the awakening of a "sleeping unicorn." Examples include artificial intelligence (the development of which underwent several "winters" before its recent flourishing), graphene, and high-power-density batteries.

A participant observed that one interesting approach to understanding the development and adoption of new technologies is to do a discourse analysis of how they are being described or associated with existing concepts—for example, a "Terminator" scenario for artificial intelligence, “Frankenfoods” for genetically engineered crops—whether by their proponents or their critics.

Inferring and Signaling Intent: Openness Versus Classification

Inferring the intent behind a technological development program may be even more challenging than detecting the new development in the first place. In fact, it is a fallacy to assume that technology development is linear, starting with intent and then proceeding through various stages of amassing capability. Sometimes undirected research, or the development of some technical capacity, creates the intent to proceed with applications. A state exploring biotechnology may not have any specific weapons motivations until it becomes more proficient with the technology, at which point the existence of that capability might cause it to pursue weapons applications more seriously. To determine intent, it can be useful to overlay social networks on top of technical networks and determine whether those developing a technology have sufficient tacit knowledge to succeed in developing weapons capabilities.

The flip side of inferring intent is signaling it: What conclusions are drawn by others when we decide to pursue some line of effort? In some cases, we can choose between exploring certain technologies—and accepting whatever conclusions others may learn from our choice—or forgoing them. In other areas, however, we cannot afford ignorance—developing a laboratory understanding of a new capability, and preventing surprise, may be critical.

A third option is to develop technology in a classified environment in the hope of avoiding inadvertent signaling. However, that option is not without risk. Openness allows anyone to work on a problem, including those who may be unable or unwilling to work in a classified environment. Many of the most creative and most productive researchers are among them. When asked in the 1980s whether more biodefense work should be classified, U.S. biodefense researcher David Huxsoll stated he would prefer to keep it open to get the right people working on it and to develop the best countermeasures. Similarly, nuclear weapon pioneer Edward Teller long argued that classification of the U.S. nuclear weapons program was holding it back, and others argued it was a fool’s errand to believe that the United States could maintain an intellectual monopoly over nuclear technology anyway.[4]

Moreover, lack of publication in some area can be an indicator that classified activity is underway in that area. Of course, it can also indicate that nothing is going on at all.

While these risks to classifying research are an important consideration, openness in defense-relevant technologies also poses risk. Published research relevant to synthesizing viruses 50 times more lethal than influenza provides information that can attract bad actors. Al Qaeda leader Ayman al-Zawahiri, in a captured internal memo, stated that “despite their extreme danger, we only became aware of [the potential of chemical and biological weapons] . . . when the enemy drew our attention to them by repeatedly expressing concerns that they can be produced simply.”[5]

Notwithstanding the risks of openness, societal vulnerabilities sometimes may need to be publicized to generate political support for the resources needed to counter them. Democracies have a hard time making substantial investments to protect against problems they do not know they have. The question is: Who is more likely to respond to such warning—those seeking to capitalize on the threat or those seeking to counter it?

Ethical and Social Components of Technology Development

Values and ethics surrounding technology development differ from culture to culture. Cultural differences, for example, have led Europe to be more reluctant to accept genetically modified crops than the United States and other agricultural exporters. Similarly, Europeans are more averse to the potential implications of data mining and artificial intelligence and have greater concerns about data privacy. The United States values predictability in operation more than many other countries or nonstate actors do—meaning that China may lag the United States in developing technology but still may be first to field it. There is value in promulgating our thinking on this—the two countries have a mutual interest in avoiding risky behaviors.

Both Chinese and American researchers discuss ethical artificial intelligence, but the Chinese literature does not stay within the Western ethical framework—the countries have different priorities. For example, U.S. guidance documents prioritize ensuring privacy against government intrusion, whereas Chinese guidance is more concerned about protecting privacy from corporations.

Prioritizing Investments in Responding to Technological Development

The plethora of potential technological threats requires that responses to them be prioritized. Infinite vulnerabilities exist, and we must decide which ones are the most realistic or most concerning and focus our limited resources on those. Technologies are not developed and deployed by individuals in isolation; they arise from a greater scientific and social community that offers multiple ways of performing analysis or staging interventions. Sociotechnical analysis can help place potential threats along a spectrum and facilitate prioritization.

Vulnerabilities to be addressed by sociotechnical analysis include both those resulting from adversaries' exploitation of new technology and those resulting from adversaries' ability to frustrate or subvert our own exploitation of new technology. (For example, the Taliban have reportedly been able to acquire U.S. military biometric devices, which may help them identify people who assisted the U.S. military in Afghanistan.[6]) As systems embodying new technologies become essential, denial of our own use of them results in damage that is proportional to the degree to which we have become dependent on them. Furthermore, penetration of these systems may not only provide adversaries with a new capability but also enable them to leverage supporting databases and resources.

We are setting ourselves up for the second category of vulnerability by paying insufficient attention to the ways information and expertise associated with the deployment of new technology travel, and by failing to accompany that deployment with rules and regulations that can mitigate unintended consequences.

We must also guard against downplaying less high-profile or low-tech threats in favor of the flashier or more novel ones involving advanced technology. Unfortunately, the introduction of new threats—such as engineered pathogens—does not necessarily push old ones—such as wild-type anthrax—off the table.

Suggestions for Further Work

Workshop participants suggested that different governance structures that might be applied to an emerging technology could be explored through a wargaming approach: What might the unintended consequences of different approaches be? Both the policy community and the technical community would need to be engaged.

Case studies would also be instructive. When technologies have been developed in the past, what consequences were feared, and which of those materialized? What were the various factors that led to the outcomes? By understanding how certain emerging technologies failed (or succeeded) to become emerging threats, we could have a better understanding of threat dynamics. Failures may teach more than successes.

Conclusion: What Makes an Emerging Technology an Emerging Threat?

What makes an emerging technology an emerging threat is multifaceted and highly dependent on context. Some workshop discussion addressed emerging *disruptive* technologies, asking whether they have the potential to radically disrupt "kill chains"—that is, the processes by which military targets are identified and forces are directed to attack them—or to radically create them. On the other hand, *disruptive* can be a positive concept in the economic context, where it refers to the potential to transform commerce and industry—recognizing that powerful beneficial applications can have a dual-use flip side.

Developing technologies can reach tipping points, at which they take on a different nature. For example, potential applications may have long been envisioned, but in practice they could not be realized without the development of additional capabilities. Lithium batteries, for example, have enabled the practical realization of many electric and electronic devices. Yet they, in turn, did not instantaneously appear, but rather evolved through their own processes and the growth of supporting infrastructure. Alternatively, some event—such as the Three Mile Island nuclear reactor accident—may trigger a societal reevaluation of a technology, markedly changing its evolution.

Many factors besides being disruptive or nondisruptive are relevant in whether emerging technologies are emerging threats, including their funding, their reliance on tacit knowledge, their institutional support/infrastructure, their use, and controls or governance approaches over them. This broad range of factors offers potential for embedding an emerging technology within a framework or governance structure that can capitalize on its positive potential and mitigate misuse.

Notes

[1] National Academies of Sciences, Engineering, and Medicine, *Biodefense in the Age of Synthetic Biology* (Washington, DC: The National Academies Press, 2018), https://doi.org/10.17226/24890.

[2] Kathleen M. Vogel and Sonia Ben Ouagrham-Gormley, "China's Biomedical Data Hacking Threat: Applying Big Data Isn't as Easy as It Seems," *Texas National Security Review* 5, no. 3 (Summer 2022), https://tnsr.org/2022/04/chinas-biomedical-data-hacking-threat-applying-big-data-isnt-as-easy-as-it-seems/.

[3] Jeronimo Cello, Aniko V. Paul, and Eckard Wimmer, "Chemical Synthesis of Poliovirus cDNA: Generation of Infectious Virus in the Absence of Natural Template," *Science* 297, no. 5583 (July 11, 2002), 1016–1018.

[4] Robert Gilpin, *American Scientists and Nuclear Weapons Policy* (Princeton, NJ: Princeton University Press, 2016), 35.

[5] Alan Cullison, "Inside Al-Qaeda's Hard Drive," *The Atlantic,* September 2004.

[6] Ken Klippenstein and Sara Sirota, "The Taliban Have Seized U.S. Military Biometrics Devices," *The Intercept*, August 17, 2021, https://theintercept.com/2021/08/17/afghanistan-taliban-military-biometrics/.

Selected Bibliography

The following bibliography provides material that was drawn on, or that supplements, ideas that were discussed during the workshop.

Abraham, Itty. *The Making of the Indian Atomic Bomb: Science, Secrecy, and the Postcolonial State*. London: Zed Books, 1998.

Abraham, Itty. "'Who's Next?' Nuclear Ambivalence and the Contradictions of Non-Proliferation Policy." *Economic and Political Weekly* 45, no. 43 (October 23–29, 2010), 48–56.

Ben Ouagrham-Gormley, Sonia. *Barriers to Bioweapons: The Challenges of Expertise and Organization for Weapons Development*. Ithaca, NY: Cornell University Press, 2014.

Ben Ouagrham-Gormley, Sonia, and Kathleen M. Vogel. "The Social Context Shaping Bioweapons (Non)Proliferation." *Biosecurity and Bioterrorism* 8, no. 1 (March 2010), 9–24.

Bresnahan, Timothy F., and M. Trajtenberg. "General Purpose Technologies 'Engines of Growth'?" *Journal of Econometrics* 65, no. 1 (January 1995), 83–108.

Cortada, James W. *The Digital Flood: The Diffusion of Information Technology Across the U.S., Europe, and Asia*. New York: Oxford University Press, 2012.

Dennis, Michael Aaron. "Tacit Knowledge as a Factor in the Proliferation of WMD: The Example of Nuclear Weapons." *Studies in Intelligence* 57, no. 3 (September 2013).

Etzkowitz, Henry, and Loet Leydesdorff. "The Dynamics of Innovation: From National Systems and 'Mode 2' to a Triple Helix of University–Industry–Government Relations." *Research Policy* 29, no. 2 (February 2000), 109–123.

Gordin, Michael D. *Red Cloud at Dawn: Truman, Stalin, and the End of the Atomic Monopoly*. New York: Farrar, Straus and Giroux, 2009.

Hymans, Jacques E.C. *Achieving Nuclear Ambitions: Scientists, Politicians, and Proliferation*. Cambridge: Cambridge University Press, 2012.

McNamara, Laura Agnes. "Ways of Knowing About Weapons: The Cold War's End at the Los Alamos National Laboratory." Ph.D. diss., University of New Mexico, May 2001.

Miller, Flagg. *The Audacious Ascetic: What the Bin Laden Tapes Reveal About Al-Qa'ida*. London: Hurst, 2015.

Montgomery, Alexander H. *Double or Nothing? The Effects of the Diffusion of Dual-Use Enabling Technologies on Strategic Stability*. College Park, MD: Center for International and Security Studies at Maryland, July 2020.

Nightingale, Paul, and Paul Martin. "The Myth of the Biotech Revolution." *Trends in Biotechnology* 22, no. 11 (November 2004), 564–569.

Rapoport, David C. "Fear and Trembling: Terrorism in Three Religious Traditions." *The American Political Science Review* 78, no. 3 (September 1984).

Silva, Patricia. "Davis' Technology Acceptance Model (TAM)." In *Information Seeking Behavior and Technology Adoption: Theories and Trends*. Edited by Mohammed Nasser Al-Suqri and Ali Saif Al-Alfi. Hershey, PA: IGI Global, 2015.

Sims, Benjamin. "Revisiting the Uninvention Hypothesis: A Transactional View of Tacit Knowledge in Nuclear Weapons Design." Paper presented at the Society for Social Studies of Science Annual Meeting, Montreal, Canada, October 10–13, 2007.

Vogel, Kathleen M. *Phantom Menace or Looming Danger? A New Framework for Assessing Bioweapons Threats*. Baltimore, MD: Johns Hopkins University Press, 2013.

Workshop Participants

Diane DiEuliis	National Defense University
Gerald Epstein	National Defense University
Amy J. Nelson	The Brookings Institution
Marijn Hoijtink	University of Antwerp
Micah Lowenthal	National Academies of Sciences, Engineering, and Medicine
Jason Matheny	RAND Corporation
Brendan G. Melley	National Defense University
Emelia Probasco	Center for Security and Emerging Technology, Georgetown University
Ted Plasse	Defense Threat Reduction Agency
David Sweeney	Defense Threat Reduction Agency
Kathleen M. Vogel	Arizona State University
Quinn Urich	National Defense University
Ryan J. Zelnio	National Science Foundation

The mission of the Center for the Study of Weapons of Mass Destruction is to prepare U.S. national security leaders to address the challenges posed by weapons of mass destruction through its education, research, and outreach programs.

CENTER FOR THE STUDY OF WEAPONS OF MASS DESTRUCTION

Mr. Brendan Melley
Director

The Proceedings series presents key discussions, ideas, and conclusions from National Defense University symposia, workshops, strategic exercises and other research, and occasionally those of international counterparts. The opinions, conclusions, and recommendations expressed or implied within are those of the contributors and do not necessarily reflect the views of the Department of Defense or any other agency of the Federal Government. For information on NDU Press, visit www.ndu.edu/press.

NDU Press