ISSUES
IN SCIENCE AND TECHNOLOGY

NATIONAL ACADEMIES OF SCIENCES,
ENGINEERING, AND MEDICINE
THE UNIVERSITY OF TEXAS AT DALLAS
ARIZONA STATE UNIVERSITY
FALL 2016

The Criminalization of Immigration

The Middle-Skill Workforce: Identifying, Training, and Employing Skilled Technical Workers

Making Sense of the New Toxic Substances Control Act

Biosecurity Governance for the Real World

The Benefits of Technocracy in China

USA $12.00 / CAN $13.00
6 3>
0 56698 57433 0

SAM WEISS EVANS

Biosecurity Governance for the Real World

Current efforts to limit the dissemination of dual-use biological research results are rooted in simplistic understandings of how such knowledge becomes dangerous. It's time for a new approach.

In 2011, controversy erupted over the publication of two peer-reviewed papers that discussed how to convert a deadly strain of influenza in birds (H5N1) to one that could be spread among mammals. How was this research allowed to happen? Was it actually dangerous? Much attention has been given to answering these questions and to considering how we might better address security concerns around biotechnology research more generally.

The H5N1 papers certainly do not mark the first time that worries have been raised about how to prevent scientific research from threatening our security. In the 10 years leading up to these publications, for example, no fewer than four reports came out of the National Research Council about science and security. Those reports, in turn, built on other efforts to govern security concerns in science since the world wars. These efforts continue today as yet another National Academies' committee, on "Dual Use Research of Concern: Options for Limited Communication," considers recommendations to control sensitive pieces of scientific studies, with attendant new procedures for how those sensitive findings would be stored and for vetting who would have access to them.

But before creating another—likely very costly—system to govern security concerns, we ought to consider the validity of the foundation on which this whole form of governance is based. Our current system builds on assumptions about the structure of knowledge and the relationship between science and the state that do not match up with actual practice. By continuing to be guided by such assumptions, we risk overplaying some concerns and expending unnecessary resources and anxiety for marginal returns on security, while missing an opportunity to enhance security in other ways.

Scientific society

Most scientists know that myriad factors—from previous research to funding sources to the pressures of the academic merit system and a drive to profit from new ideas—have a significant impact on the content and direction of research programs. Current research sits within an ecosystem of sponsors, universities, federal and state oversight, and other factors that principal investigators must quickly master if they are to advance their careers. Similarly, scientists must have at least an unconscious appreciation of the importance of things such as lab training to build the skill set necessary to understand the knowledge transmitted in publications and to use those skills and training to create new knowledge and applications. When a paper says, for instance, that a certain method was used, unless a scientist is familiar with that method, the ability to assess and use the knowledge is limited.

Yet when it comes time to consider how best to govern security concerns around science, scientists and security experts often rely on radically oversimplified assumptions about scientific knowledge and its relation to the state and society. These assumptions are familiar to many people, because they are the bedrock of most science education: science is the pursuit of truth unburdened by any concerns about society; if we can just get politics and corporate interests out of the way, scientists will uncover the truths of the universe, and we will all reap the rewards; and science advances through the accumulation of discrete facts and causal assertions—the world is round, nothing goes faster than the speed of light, pathogen A causes disease X.

But basing our governance system on these assumptions is dangerous because they eliminate the social aspects of science, and thereby most of the context within which a governance system must function. Yet this is exactly what we have done, from controlling military research in the world wars to current attempts to govern open biological research that has no military ties.

Governing knowledge

Although there are a few examples of states controlling scientific knowledge before the world wars, the current system for governing security concerns in science has two primary starting points: the Atomic Energy Act (AEA) of 1946 and President Reagan's 1985 National Security Decision Directive 189 (NSDD-189). The first of these established the idea that limiting access to certain facts and findings through classification was an acceptable mechanism for controlling scientific information that had a national security concern. Unlike wartime-specific controls on science, mostly notably those around the Manhattan Project, the AEA was the first legislation to restrict the flow of scientific information with no real end date envisioned for the controls. Section 10 of the act states that knowledge about nuclear weapons, nuclear energy, and related production processes is automatically classified by the state regardless of where it was produced.

If the AEA established that classification was an acceptable mechanism to govern security concerns in science, NSDD-189 solidified the idea that classification was the *only* acceptable mechanism. Section III of NSDD-189 says, "It is the policy of this administration that, to the maximum extent possible, the products of fundamental research remain unrestricted. It is also the policy of this administration that, where the national security requires control, the mechanism for control of information generated during federally-funded fundamental research in science, technology, and engineering at colleges, universities, and laboratories is classification."

Both of these documents have a common assumption about the shape of knowledge: it exists in discrete and recognizable chunks that can be allowed to flow or be withheld at a state's discretion. This way of structuring knowledge leads to governance systems built around lists, such as the Select Agent List or export control lists. Assuming potentially dangerous knowledge comes in discrete chunks is acceptable if those involved in classifying the knowledge can clearly define which chunks are worrisome, can know from whom the knowledge should be kept, and can actually keep the knowledge from flowing to the malicious user. It is much easier to tell which chunks are worrisome if that knowledge is produced within a security environment such as a weapons lab. In biology, where much of the research is conducted in unclassified and decentralized settings, determining danger is much more difficult. Similarly, in an age of non-state-sponsored violence, differentiating friend from foe at the individual or even lab level takes substantially more resources and intelligence gathering than when conflict was mostly between nations. Such complexities require an appreciation of the factors within a research environment that would make a piece of knowledge dangerous.

The genesis of NSDD-189 also shows how the idea that dangerous knowledge comes in discrete chunks was combined with the assumption that fundamental science is autonomous to create the governance system we are stuck with today. In the 1970s, concern was growing among the defense community that, although export controls were used to slow the process of potentially dangerous technologies falling into nefarious hands, this governance mechanism had no ability to control the flow of ideas and knowledge, and amendments were made to add knowledge to export control lists. Yet the very idea of controlling the flow of knowledge conflicted with the belief, established in the 1950s and 1960s, that unfettered scientific research was the bedrock of progress. Indeed, in 1982, the National Academy of Sciences sought to counter the encroachment of government regulations on fundamental research when it issued the Corson report, *Scientific Communication and National Security*, which argued that the national security and economic well-being of the nation were actually *supported*, not undermined, by maintaining unfettered basic scientific research. Thus, the report argued that the vast majority of research should have "no restriction of any kind limiting access or communication," while allowing that "government-supported research [that] demonstrably will lead to military products in a short time" could be subject to classification. For the gray areas between these two categories, it recommended only that foreign nationals not directly participate in the research and that the government have a right to see articles before publication, but not to modify them. NSDD-189 was, in part, a statement of political support for the Corson report and was a reaffirmation that science best serves society when it is allowed to go where it pleases.

Although this formal assertion of scientific autonomy helped quell the debate between scientists and the national security apparatus, it doesn't mean that this vision of autonomous science matches up with reality. Although many universities do not conduct classified research, for instance, a substantial amount of funding has come from commercial and defense sources, sometimes with significant strings attached. The Department of Defense (DoD) provided an immense investment in scientific research during the Cold War on everything from transistors to environmental monitoring, and its goals have strongly influenced the direction and content of scientific research. In fact, political systems, institutional bureaucracies (or lack thereof), ethical persuasions, and guiding visions of the type of society we hope to live in are all ingrained in decisions about what scientific knowledge we produce. This is intimately understood by the scientific communities that have orchestrated top level policy attention to their work, such as the Human Genome Project, the National Nanotechnology Initiative, and the BRAIN Initiative.

Does DURC work?

Scientists, as a matter of principle, call foul at any attempt at intervention but in practice understand and shape the societal context within which their work can flourish. The latest iteration of our security governance system is the concept of "dual-use research of concern" (DURC), coined in 2007 by the National Science Advisory Board for Biosecurity (NSABB), a federal advisory committee managed by the National Institutes of Health. DURC takes as its starting point the claim that much knowledge in the life sciences could be maliciously applied, but argues that only a significantly smaller set of it is dangerous and warrants extra attention. DURC is now defined as federally funded "life sciences research that, based on current understanding, can be reasonably anticipated to provide knowledge, information, products, or technologies that could be directly misapplied to pose a significant threat with broad potential consequences to public health and safety, agricultural crops and other plants, animals, the environment, materiel, or national security."

What should be done when an area of research is labelled as DURC? The H5N1 influenza case was the first major attempt to answer this question. Federal moratoriums were issued, first on H5N1 research for a year, and then on a wider swath of "gain-of-function" research, a policy that is still in place as of late 2016. A moratorium, like classification, is a very blunt tool for governing security concerns, and the scientific community accepted its use only because they understood that moratoriums would be limited in duration. The National Academies held several symposia, the NSABB produced a range of reports, and the government issued two policies on the oversight of DURC.

Yet because these policies are built on the beliefs that knowledge is made of discrete chunks and security must be balanced against academic freedom, they are of limited value in protecting security. DURC oversight covers only US federally funded research—neither philanthropic nor corporate research is covered—and only in relation to items on the Select Agents List, and even then only when one of seven types of experiments are conducted. The inclusion of these seven types of experiments was itself the result of another line-drawing attempt by the National Academies, the 2004 Fink report, *Biotechnology Research in an Age of Terrorism*, though this list was meant to be only suggestive, not definitive. In defining DURC this way, several biology research projects, such as the 1997 Penn State project to aerosolize non-select agents to bypass the natural defense of the lungs and the 2001 de novo synthesis of the polio virus, which had prompted the development of the DURC concept in the first place, were still not covered. Current research on gene drives, which could be designed to alter or even kill off entire wild populations of organisms, are also not covered as long as they do not use any select agents. Moreover, even if the list of select agents or experiments were to be modified as, for example, the pathogenicity of novel synthetic organisms becomes known, that knowledge may not be developed until well after the organism has been characterized extensively in the open literature, making post-hoc control of information virtually impossible.

DURC oversight is also limited because most journals and universities lack appropriate staff with access to information about possible ways the knowledge might be misused, or with the training to conduct a risk assessment if they did have that knowledge, so DURC research is likely to go unrecognized, let alone adequately governed. In large part, this is due to the narrow focus of any oversight committees a university might have and the lack within scientific training of discussion of how research might cause harm._

These limits and contradictions are an inevitable outgrowth of policy aspirations founded on a simplistic picture of how science operates, and they compromise the nation's ability to understand and govern security threats that emerge from the growth of scientific knowledge.

Reducing the threat

What if, instead, the United States were to craft a security governance system based on what is actually known about how knowledge is produced and disseminated? Several insights would immediately become relevant. Scientists' training and resources heavily influence what topics they research and how they conduct their work. By analyzing this training and the environment within which research is conducted, we can gain a much richer understanding of the context within which the research they conduct and publish might merit concern. An analysis of the context of research also will provide a better understanding of how a lab makes use of tacit knowledge that is not formally presented in publications but rather is passed on when scientists work together on experiments and idea development.

The potential value of such an approach is well illustrated by the H5N1 case, where initial concerns within the NSABB about the two papers were alleviated when the authors added discussions of the wider context of epidemiology research within which the studies occurred, the safety and security measures the researchers took to protect themselves and the public, and the goals and public health benefits of the research. The additional information allowed for a security assessment that satisfied the NSABB that the descriptions of the experiment did not constitute a threat. The solution wasn't to communicate less, but to communicate more!

Knowledge, here, is a fluid entity, and is more about the connections among ideas, people, and the environment in which it is produced than it is a set of discrete facts produced by a single individual in an ivory tower. Of the infinite ways scientists might write a paper about a finding, which knowledge they make explicit and which they leave tacit depends on how they were trained and what rewards and punishments might come from saying—or not saying—certain things. What does this understanding of knowledge production suggest about how we might recognize and govern science-related security threats?

First, those involved with guiding policies must recognize that directions of knowledge production reflect choices made within political, ethical, and institutional contexts, and within such contexts we already limit in many ways the types of knowledge we find it acceptable to produce. Scientific publications are heavily structured by what the people producing and reading them think is acceptable to include. There is an infinite number of details that are not included in these documents because they are seen as extraneous or taboo within the research culture.

Scientists should be better rewarded for openly reflecting on the ethical choices and the safety and security environment of their research, in particular on how their vision of how their research might benefit society and the environment might be someone else's vision of harm. A simple step toward this change in reward structure for the life sciences might be taken from computer science, where security professionals—"white hats"—specialize in testing computing systems to ensure their security, explicitly showing the limitations of particular configurations of code. A similar community should be built in the life sciences to work closely with those producing new knowledge.

It is not just the scientists who need to become more reflective about the contexts of their work. The security community has its own limitations on how it produces threat assessments based on the assumption of science as a discrete set of facts. As Kathleen Vogel of North Carolina State University said in her analysis of intelligence agencies' initial conclusion that the H5N1 papers were a security risk, "with the proper training in science and ethnographic methods, one intelligence analyst or contractor working over a ten-day period could have gathered new, substantial information about the H5N1 experiments from site visits ... [to laboratories that]would have yielded a wealth of new information about the experimental work that was not available merely by reading the manuscripts."

The H5N1 case helps make clear that governing security concerns raised by science is not just a matter of controlling the transmission of facts but of understanding how the goals of states, companies, and citizens shape, and are shaped by, decisions about how to direct research and innovation. If scientists want to do something malicious with a piece of knowledge, they will need much more than the published article. Monitoring, therefore, needs to attend not to any particular piece of knowledge, but to the broad array of factors (training, resources, tacit knowledge, intent to do harm, and so on) that combine with things such as scientific publications to produce a credible threat.

Second, rather than building fences around narrow objects of concern, we should be building conversations across areas of relevant expertise. A promising example of such an effort is the Federal Bureau of Investigation's Safeguarding Science Initiative within the Weapons of Mass Destruction Directorate's Biological Countermeasures Unit. This unit focuses on building long-term relationships with the broad range of labs and researchers who are conducting, or might conduct, research with potential security concerns. Instead of creating an antagonistic relationship with the scientists, the initiative is focusing on becoming the place that sensible scientists would turn to when they have a question about the security aspects of their work. This approach recognizes that both scientists and security professionals have areas of expertise that need to be meshed together to understand what should count as research of security concern. It also places an emphasis on monitoring developments—in knowledge, but also in resources, intent, networks, and so on—in addition to creating static lists of objects of concern and particular points in the knowledge production and innovation cycles where assessments occur. The Safeguarding Science Initiative should be studied for its strengths, weaknesses, and applicability across subfields of life sciences that raise potential security concerns.

Third, security itself should not be considered in isolation from the broad range of values that motivate the quest for knowledge. Once we start to appreciate that it is impossible to separate the production of knowledge from the societal goals that guide and are supported by that knowledge, the debate can shift from whether a government should or should not intervene in science to a discussion about what types of societal goals we want to incorporate in the research we promote. From this perspective, security concerns may even prove complementary with a wider set of other concerns that a society may have. For example, within the world of export controls, countries work closely with companies that make military and dual-use technologies because the goals of intellectual property protection and national security often align when trying to control the flow of information that is sensitive for both economic and security reasons.

As the National Academies and the US government continue their efforts to govern the security concerns of biological research, they should move away from the limits of the DURC concept. The first step down this path is to stop using the polarizing discourse of security emergencies and academic autonomy. Knowledge is always produced within a social context, and security is only one of many goals a society is striving for. From that starting place, we can begin to build a contextual approach to governance that is appropriate for the complex practice of real-world science.

Sam Weiss Evans is a visiting research fellow with the Program for Science, Technology and Society at the Harvard Kennedy School and a research affiliate at the Program on Emerging Technologies at the Massachusetts Institute of Technology.

Recommended reading

Philip Campbell, "Issues in Dual-Use Biology Publishing," Presented at the National Academies of Sciences, Engineering, and Medicine's Committee on Dual Use Research of Concern: Options for Limited Communication, 1st Meeting, Eventi Hotel, New York, NY (July 11, 2016), available online: http://sites.nationalacademies.org/cs/groups/pgasite/documents/webpage/pga_173610.pdf.

National Research Council, *Science and Security in a Post 9/11 World: A Report Based on Regional Discussions Between the Science and Security Communities* (Washington, DC: National Academies Press, 2007).

_____, *Understanding Biosecurity: Protecting against the Misuse of Science in Today's World* (Washington, DC: National Academies Press, 2010).

_____, *Research in the Life Sciences with Dual Use Potential: An International Faculty Development Project on Education about the Responsible Conduct of Science* (Washington, DC: National Academies Press, 2012).

_____, *Perspectives on Research with H5N1 Avian Influenza: Scientific Inquiry, Communication, Controversy: Summary of a Workshop* (Washington, DC: The National Academies Press, 2013).

_____, *Potential Risks and Benefits of Gain-of-Function Research: Summary of a Workshop* (Washington, DC: National Academies Press, 2015).

Kathleen M. Vogel, "Expert Knowledge in Intelligence Assessments: Bird Flu and Bioterrorism," *International Security* 38, no. 3 (2014): 39-71.