

When All Research Is Dual Use

Governing new biosecurity threats is not merely a matter of good intentions and better training; it requires a paying proper attention to the social contexts of science.

In a March 2022 paper in *Nature Machine Intelligence*, researchers from a US pharmaceutical company who were building artificial intelligence systems for virtual drug discovery issued a wake-up call to their colleagues. After years of working on a suite of models to improve toxicity prediction, the researchers were invited to an international security conference to give a presentation on how such models could be misused to create chemical and biological weapons—something they had not previously considered, even though they had worked with neurotoxins and Ebola. “The thought had never previously struck us,” they wrote. By simply changing their models to search for molecules with more toxicity rather than less and running the trained algorithm for under six hours, the researchers were able to generate 40,000 molecules that were likely lethal, including the nerve agent VX and many new molecules that were predicted to be even more potent than known chemical warfare agents. “We were naive in thinking about the potential misuse of our trade,” the researchers wrote. “We are not trained to consider it.”

For biosecurity, the future depends, in part, on ensuring that wake-up calls occur at the beginning of the design cycle, not after there is already a product and the developers just happen to be invited to a conference with a security focus. But fixing a situation of this nature—where scientists are neither trained nor

rewarded for attending to the societal consequences of their research and rely on serendipity to understand new threats—is not merely a matter of better education and oversight.

The issue is much larger. The debate about the relationship between science, security, and society has reached a new crossroads. For most of the last century, the science and policy community has chosen a path that built a scientific-industrial structure based on the assumption that the best way to maximize the societal benefits from science was to leave it to its own devices, divorced from the society it ostensibly serves. But as the world grows ever more complex, is this still the best way? Was it ever?

The problems with the myth of asocial science, and its accompanying pantheon of lone hero scientists, are widespread and well known—but not, it seems, to policymakers, who continually reinscribe it. The myth can be found throughout US research, innovation, and governance systems, all of which fail to incentivize scientists to engage with society—or, often, even with those from other fields of study who might bring a different perspective.

What is needed is the opposite: recognition that science is a social system. In this complex social system, what questions get asked—like whether an AI tool could be used for ill—has as much to do with institutional culture, economics, politics, and ethics as

with the science itself. The knowledge and technologies created by this system are a result of the contexts where they are made, and will bring about new ways of harming as well as helping.

The United States has a moment now to construct a new relationship between science, security, and society, at least within the life sciences. The National Security Commission for Emerging Biotechnology, established through Section 1091 of the 2022 National Defense Authorization Act, will soon begin its deliberations. According to the legislation, this commission, modeled on the earlier National Security Commission on Artificial Intelligence, has a broad mandate to “consider the methods, means, and investments necessary to advance and secure the development of biotechnology, biomanufacturing, and associated technologies by the United States to comprehensively address the national security and defense needs of the United States.”

The commission’s recommendations for modifying

the misappropriation of research and development to the detriment of national or economic security, related violations of research integrity, and foreign government interference.”

This framing, however, reinscribes a “fortress America” mentality that has been roundly criticized by the National Research Council as “quietly undermin[ing] our national security and our national economic well-being.” Guards, gates, and guns only help when it’s clear what the threats are and what is to be protected. In the world of emerging biotechnology, neither is clear. More attention must be paid to how current forms of security governance undermine the very security policymakers are trying to achieve.

Deliberations on the content of research present a moment for the commission to choose a new path. In the early 2000s, several high-profile publications and the 2001 anthrax attacks sparked renewed attention to the relationship between biology, security, and society,

“We were naive in thinking about the potential misuse of our trade,” the researchers wrote. “We are not trained to consider it.”

the governance of biology, and biosecurity in particular, should stem from an understanding of science as a social system. Training, funding, research, publication, innovation, and oversight all need to be radically altered if decisionmakers are to put the social context of biological research and technology at the heart of policy. At each turn of the negotiations, the easy answer will be to return to the myth of science as asocial, but doing so will only undermine the future well-being of our society.

All research is dual use research

Before diving into the details of how biotechnology research should be done, the commission should start by looking hard at the assumptions embedded within some basic definitions. “Research security” and “research integrity” are concepts that build on the belief that, while science is by and large inherently good, certain actors may be bad, and ensuring the proper conduct of research means separating the good actors from the bad ones. For example, the recent guidance on implementing National Security Presidential Memorandum 33, the policy regarding national security issues around government-supported research and development, defines research security as “safeguarding the research enterprise against

providing several opportunities to reform biosecurity governance around an understanding of science as a social process. Instead, the government drew hard boundaries, developing policies on dual use research of concern and governance guidance on pathogens with pandemic potential, limiting concern to pathogens and “knowledge, information, products, or technologies that could be directly misapplied to pose a significant threat.” This focus overlooked many ways biological innovation could become concerning—for example, AI-bio convergences or organisms modified to cause other types of harm than disease. Over time, the potential for such diverse, unknown threats has grown, but current oversight puts substantial resources toward research that officials know might be concerning, with barely any left for the rest. The most likely places for new security concerns to arise are precisely those areas of science that have the least security attention, and that is because of the assumption that good intentions lead to good outcomes.

This dynamic must change, and it is the commission’s responsibility to provide an example of how that could be done. The commission should expand its understanding of biosecurity, embracing the point that *all research is inherently dual use*. In the past, security governance has

been compliance-based, focused on lists of known pathogens and known entities of concern. This pointillist approach will have very limited, though still essential, utility moving forward. Having special governance only for known threats makes little sense when new security concerns are likely to emerge from the unknown—whether from unintended consequences, natural evolution, or malicious use. What is needed instead is a curiosity-based system that attunes researchers, funders, and policymakers to attend to security throughout the research lifecycle.

Starting from an understanding of all research as inherently dual use and identifying these unknowns as early as possible requires a shift: moving away from the belief that good *intentions* make good research and moving toward paying *attention* to whose security and liberty is privileged and whose is undermined in the governance system. The researchers at the start of this article had every intent to do good, but were trapped by

adaptive governance styles. Security governance needs to go beyond encouraging compliance with best practices for known concerns; it must account for threats where the objects of concern are uncertain, the actors more diffuse, and the responsibilities more distributed. Taking seriously the social fabric of science demands including security as another of the threads, but it is a thread that needs to be woven throughout.

To accomplish this, security must become a forethought to innovation, not an afterthought. Training students in how to understand, question, and reform the social aspects of science and technology needs to be part of science education from high school through postgraduate programs. This training should alert young scientists to the ways systems of governance ensure some types of security better than other types. It should empower them to think about security and safety more broadly, and to incorporate that thinking into their own work. The International Genetically Engineered Machine

The most likely places for new security concerns to arise are precisely those areas of science that have the least security attention.

a system that has balanced the liberty of the academy and economy with national security by choosing to consider only a very limited subset of knowledge and technology as dual use.

Paying attention to questions of whose security and liberty matter will require networks of expertise spanning the natural, social, and economic sciences working with biosafety, intelligence, and security professionals to arrive at much richer understandings of what could constitute a threat, and to agree on what acceptable governance looks like. An essential part of this process will be including voices that have historically been missing, especially from populations who have benefited less from previous research and often bear the brunt of negative consequences.

From intention to attention: the need for diverse, experimental governance

To bring attention to the social nature of science, the commission could champion a more experimental approach to governance that focuses on sandboxing, documenting, and learning from governance experiments. To begin, the commission could build a capacity for testing anticipatory, participatory, and

(better known as iGEM) competition is a good example: every year, 6,000 synthetic biology students from around the world are trained in how to put the social context of biology first in their project designs.

A key part of reform will be encouraging the creation of better forums to debate what constitutes a security concern, whose security matters, and how security can be achieved. Of course, such reforms are likely to face a backlash from entrenched institutional structures. Most researchers lack any incentive, much less the knowledge and resources, to engage in such conversations—another lingering symptom of science’s separation from society. But this type of activity needs to become part of all researchers’ jobs. Getting there requires reevaluating a system that rewards “research productivity” by measures that exclude attention to how that research shapes, and is shaped by, social factors. The commission, with its national stature and congressional charter, is well positioned to change this dynamic, and to catalyze community involvement and investment in creating a secure biological future.

One area for focused improvement involves federal research funders. High-level policies should explicitly state the need for integrated funding that encourages

attention to the broader aspects of research. The “broader impacts” criterion used by the National Science Foundation (NSF), for example, could be enhanced to explicitly include activities that experiment with weaving security governance throughout the research lifecycle. The commission could build on current experiments in funding redesign, such as the Defense Advanced Research Projects Agency’s advisory groups for the legal, ethical, environmental, dual use, and responsible innovation aspects of its programs. Formed by program managers as they build their programs, these advisory groups have impacted the construction of funding calls as well as the conduct of research that is funded by the program. Federal funders should be encouraged to identify, try, and adapt new models, then reflect on their own experiments. Framing these efforts to change funding structures as experiments themselves highlights the need to document assumptions, share what works and doesn’t,

The commission itself can also be understood as a form of experimentation in security governance. One of the bedrock beliefs it should investigate is that the tension between academic freedom and national security is still adequately resolved by the 1985 Presidential National Security Decision Directive 189 (NSDD-189), which states that “fundamental research” should not be subject to any form of security oversight unless it is classified. NSDD-189 was the result of a particular Cold War moment in the debate between science and security; the actors, content, and context of science have changed dramatically since the 1980s.

While NSDD-189 was a bargain between the government and the scientific enterprise, what is needed today is a bargain that also includes civil society and industry. This is not an invitation for more regulations and reporting. It is an invitation to a forum for continuing dialogue about the relationship between science, security, and the state. What many observers

Having special governance only for known threats makes little sense when new security concerns are likely to emerge from the unknown.

and analyze the results systematically and openly.

Another aspect deserving high-level attention is peer review. Although the process is a bedrock of science, to better interweave science and society, who constitutes a “peer” needs to be reconsidered. NSF and other funding agencies could deliberately diversify their grant review panels to include researchers from disparate fields, and even practitioners from beyond the academy, who could evaluate proposals in terms of their societal aspects. The Department of Energy’s Joint Genome Institute is already experimenting with this by including security, social science, ethical, and legal experts alongside scientists in its review process. Some scientific journals are experimenting with security review of manuscripts, including assessing potential concerns by using the Materials Design Analysis Reporting Framework, which names best practices in transparent reporting of study design. Expanding the types of expertise that are considered peers to include social scientists, practitioners, and, when relevant, members of the intelligence and other security communities, will require delicate footwork by the commission, as it treads very close to questions of the autonomy of the research community.

forget is that NSDD-189 has lasted for as long as it has because of the equal footing that academia and the government had in the negotiations to develop it. The commission would be wise to recreate that dynamic, both in its conduct and in the recommendations it develops.

Putting the social aspects of science at the heart of the commission’s work will not be easy, as it questions some of the underlying assumptions of science—and of national security—for the last century. But the world in which those foundations were laid down no longer exists. In moving from intention to attention, experimenting with governance mechanisms, and bringing a wide range of voices into the room, researchers reduce their reliance on both serendipity and the guards, gates, and guns approach that is no longer sufficient to protect society. Embracing the social aspects of biology will allow the science and policy community to realign security and liberty for the twenty-first century, and—best of all—it will do so by drawing on America’s strengths as a democratic nation.

Sam Weiss Evans is a senior research fellow in the Program on Science, Technology, and Society at the Harvard Kennedy School.