# Title: Experimentation in Biosecurity Governance

**Authors:** Sam Weiss Evans[1,2,3*], Jacob Beal[4], Kavita Berger[5], Diederik A. Bleijs[6], Alessia Cagnetti[7], Francesca Ceroni[8,9], Gerald L. Epstein[10], Natalia Garcia-Reyero[11], David Gillum[12], Graeme Harkess[13], Nathan J. Hillson[14], Petra A.M. Hogervorst[6], Jacob L. Jordan[15], Geneviève Lacroix[16], Rebecca Moritz[17], Seán S ÓhÉigeartaigh[3], Megan J. Palmer[18], Mark van Passel[6]

**Affiliations:**

[1] Program on Science, Technology, and Society, Harvard University.

[2] Program on Emerging Technology, MIT.

[3] Centre for the Study of Existential Risk, University of Cambridge.

[4] Raytheon BBN Technologies.

[5] Gryphon Scientific.

[6] Netherlands Biosecurity Office, National Institute for Public Health and the Environment, Bilthoven, The Netherlands.

[7] Polo d'Innovazione Genomica Genetica e Biologia (PoloGGB), Terni, Italy.

[8] Department of Chemical Engineering Imperial College London.

[9] Imperial College Centre for Synthetic Biology.

[10] Center for the Study of Weapons of Mass Destruction, National Defense University.

[11] Engineer Research and Development Center, U.S. Army.

[12] Arizona State University.

[13] The Pirbright Institute.

[14] Joint Genome Institute, U.S. Department of Energy.

[15] Nuclear Threat Initiative.

[16] Centre for Biosecurity, Public Health Agency of Canada.

[17] University of Wisconsin-Madison.

[18] Center for International Security and Cooperation, Stanford University.

*Correspondence to: Samuel_Evans@harvard.edu

**Abstract:** Biological research and its applications are rapidly evolving, and their governance should as well. This piece argues for thinking of biosecurity governance itself as an experimental space. Doing so focuses attention on assumptions in the governance process and its iterative and continuously evolving nature. We draw out three lessons for better working with experiments in biosecurity governance.

**One Sentence Summary:** Thinking of biosecurity governance as itself an experiment enables a stronger focus on methods to develop systematic learning.

**Main Text:**

Biological research and its applications are rapidly evolving, and their governance should as well. Life scientists are today joined by computer scientists, engineers, material scientists, chemists, and others in manipulating biological materials for many purposes. The pace and nature of advances are straining current governance. New attempts at the governance of biology are emerging in response, rethinking traditional assumptions about how science works and who is responsible for governing. Unlike scientific experiments, however, these governance approaches are often not evaluated, analyzed, or compared. This hinders building a cumulative base of experience and opportunities for learning.

Consider "biosecurity governance," a term with no internationally agreed definition, here defined as the processes influencing and shaping behavior to prevent or deter misuse of biological science and technology. Traditional international biosecurity efforts have focused largely on risk management (i.e., addressing accidental and deliberate risks from pathogens and toxins) and dual use research (i.e., potential malicious exploitation of knowledge, skills, and technology). These efforts assume we already know what to worry about (known pathogens and toxins), and that we can control access to them, even if those who carry them out understand the shortcomings and limitations of these assumptions (*1*, *2*).

In the last decade, however, our ability to manipulate living organisms and entire genomes has rapidly advanced through the development of novel tools like CRISPR, improved sequencing techniques, and fast-paced synthetic genomics approaches (*3*). This has improved our ability to generate novel pathogens, cell types, animals, materials, and gene drives, all of which may represent additional security concerns. Moreover, there has been increasing interest in non-state actor utilization of biological weapons (*4*, *5*). These changes in technology and political environment, coupled with natural disease evolutions like COVID-19, are testing existing governance processes.

Various communities involved in governing the life sciences are now looking beyond existing biosecurity models, policies and procedures. Three brief examples show how changes, while innovative, are currently sporadic and often ad hoc responses to identifying particular security deficiencies.

First, after heated debate about two experiments modifying avian influenza in 2011, the US Government developed policies on dual use research of concern (DURC), requesting scientists initiate added research reviews under certain conditions. Assuming such added oversight will only be implemented if minimally invasive, the policies restricted oversight to a subset of work on a subset of known pathogens. Recognizing that these policies still focus on known pathogens, in 2017, the US developed an additional policy focused instead on post-experiment attributes of an organism. This Potential Pandemic Pathogen Care and Oversight (P3CO) policy is also the first to consider whether research is ethically justifiable as well as scientifically meritorious. Reviews on both of these policies have only just started (*6*), even though stakeholders have called for the DURC policy to be reviewed for years (*7*).

Second, a decade ago the Federal Bureau of Investigation's Biological Countermeasures Unit decided that preparing for potential biosecurity events required staying abreast of advances in biology and engaging closely with the life science research community, including emerging "DIY" community labs, universities, and companies (*8*). This meant building both internal scientific expertise and community liaison capacity. Both perspectives were not necessarily shared widely within the FBI, nor in the public's image of this organization and how it operates (*9*). Moreover, these moves made scientists more responsible for identifying potential security concerns, whereas previous governance measures often assumed security was not something scientists needed to have in mind.

Third, the American Biological Safety Association (ABSA) International observed that biosafety professionals have been increasingly asked to assess security as well as safety aspects of research, but do not know a) how to assess security concerns, and, perhaps more importantly, b) how to think about malicious intent and intentional release. They assumed further training would improve research security, and have actively worked to create more biosecurity educational opportunities, including development of a global biosecurity credential (*10*).

These sporadic changes are potentially very productive, but without systematic analysis and learning across them it is hard to know. To improve, we suggest the metaphor of 'experiment' opens a useful new approach to biosecurity governance. We should consider attempts at rethinking biosecurity governance as opportunities to experiment with new sets of assumptions about the relationship between biology, security, and society, leading to hypotheses to assess and iterate upon. After all, we do not have perfect knowledge of ways biology might be misused, or the best ways to prevent misuse. There is no *a priori* reason to believe our original assumptions and hypotheses are optimal. The consequences of getting assumptions wrong, such as a global synthetic pandemic, are perhaps the strongest argument for testing a wide range of assumptions in ways that can provide signals of effectiveness prior to catastrophic events.

An experimental approach focuses attention on the need to be systematic and open about analyzing the existing systems' limitations and promoting experiments addressing or working around them. It also means developing better methods to collect data to evaluate effectiveness of governance, coupled with data sharing across experiments and with future experiments. These meta-level discussions are key for any robust and adaptive biosecurity governance system (*11*).

The experimental metaphor does have some limitations. Security governance strategies are designed not to fail catastrophically, and governance has many actors involved in design and implementation. It is perhaps better to understand our use of 'experiment' here in terms of deliberate social experiments around the introduction of new technology, where the focus is on uncertainty, lack of control, and systematic learning (*12*). This places the concept closer to a design-build-test cycle, but with the focus placed not on controlling the system, but on governing in a complex adaptive space.

**Biosecurity governance as an experiment**

One current experimental governance approach is the international Genetically Engineering Machines (iGEM) Foundation's Safety and Security Program. iGEM runs a yearly competition for around 6000 students and community biolab members from over 40 countries. Each year, iGEM generates a set of hypotheses about how its proposed changes in safety and security governance might affect teams and lead to better oversight, and reviews cases that tested, or were not caught by, its system. Through these reviews, iGEM recognized that processes for screening team genetic sequences for known pathogens both provided false positives and missed work with potential security implications beyond issues with known pathogens. This in turn led to iGEM transitioning to a function-based, rather than sequence-based, screening architecture. This is all part of iGEM's commitment to a multi-tiered, iterative security program addressing an adaptively expanding range of likely concerns (*13*).

Thinking about biosecurity governance as an experiment focuses our attention on several oft-underappreciated aspects of governance. The first is the set of assumptions we make in the process of governing, most notably about the structure of science, governing authorities, and how those relate to specific conceptions of security. These assumptions tend to come in packages. Employing a system like export controls, for example, relies on an assumption that science consists of discrete knowledge entities (e.g., published articles or biological specimens), restricting the export of which leads to security. This, however, carries an assumption that threats are likely to originate abroad, as opposed to, say, assessing the insider threat within labs in a country. Another example is the assumption that scientists are best placed to govern themselves, which is at the heart of the DURC policies, even though scientists may not have necessary training to identify security concerns. This assumption is so firmly rooted in biosecurity governance that it is difficult to question it, and when questioned, to have evidence that contradicts it feed back into governance redesign (*14*). In drawing out these assumptions and comparing them across experiments, we can more systematically understand the contexts in which they are likely to hold and where experiments based on different assumptions might be more productive.

A second underappreciated aspect of governance is its iterative and evolving nature. Governance processes and the associated communities continually renew, both in response to changing technological capabilities and to changing community and societal conditions. We can take advantage of this to learn from past governance experiments. Currently, learning from governance experiments usually occurs through *ad hoc* meetings and publications originating from an organic desire to share experiences (*15*) or from a broader strategy to create space to talk about lived experiences (*16*). The co-authors were part of a recent workshop at the University of Cambridge called *Novel Practices in Biosecurity Governance*, which aimed to shift discussions from focusing on how to address individual security governance challenges to instead focusing on the value and methods of sharing and iterating across experiences relating to a range of biosecurity governance experiments.

**Learning across experiments**

Those who fund life science work, oversee it, set or carry out policy regarding it, or engage in it as a researcher, citizen, or other interested party, may want to experiment with different ways of understanding what counts as a security concern and what should be done about it. In the spirit of learning across experiments, three initial lessons came out of the Cambridge workshop.

First, in designing a governance experiment, consider framing the proposed set of actions in terms of hypotheses, which in turn are based on a set of assumptions about the science, the governing authorities, and how those relate to a specific conception of a security concern. For example, early presentations given to biotechnology-related groups by the FBI Special Agent in charge of the Biological Countermeasures Unit, Edward You, clearly assumed that biosecurity was different than nuclear or chemical security because biological pathogens already exist in the environment and could be created by people without large facilities or state sponsorship. You's proposed solution was a governing structure that mirrored this dispersed scientific environment, one that was collaborative rather than top down. You often played up the difference in approach he was trying from the image the public had of the FBI: "We are the FBI, we're with the US Government. . . and we're here to help. Really. And we need you!" (*17*). If the FBI had considered its proposed solution as a hypothesis, it could also have developed a set of metrics to be able to assess, from the beginning, whether such hypotheses held up, and if not what might need changing. This might involve, for example, structured feedback from community labs about FBI engagement, and routinized sharing across field offices of standard procedures for developing community relationships. It can be very helpful in both designing and documenting experiments to work with social scientists and anthropologists who can identify hidden assumptions and develop alternatives that might better align with the goals of governance (*18*).

Second, it is useful to have a capacity to quickly identify difficult or unanticipated cases and adapt governing processes to account for them. Use moments where the governance system encounters difficulties to reflect and update the assumptions being made about the science and its contexts. To the extent possible, sharing case studies (including both failures and "near misses") in a timely fashion could aid other biosecurity processes greatly. iGEM developed this capacity, as noted above, and quickly put it to work when a 2016 student team claimed to be developing a gene drive. After working closely with the team and experts to understand exactly what was and was not accomplished, iGEM was one of the first places to produce a policy on gene drives (*19*). It then wrote up its lessons learned and shared them with the wider biosecurity community at a workshop the following year.

Third, learning involves connecting with communities who have tried similar experiments, and who could build earlier results. These groups range from networks of community biolabs to international efforts like the Global Health Security Agenda's action package on biosafety and biosecurity. One example of how to learn across these groups is to engage with the Emerging Biosecurity Leadership Initiative (ELBI), run by the Johns Hopkins University Center for Health Security, which provides both a broad training in biosecurity and an active network of alumni to consult. There are also specific fora like the Biological and Toxin Weapons Convention Meeting of Experts, or non-state venues like the ABSA International Biosecurity Symposium scheduled to be held in May 2020. Developing communication across communities means addressing barriers to communication, such as industrial considerations of competition-sensitivity, governmental controls (e.g., export restriction, classification), differing terminology, and the time and resource costs imposed by in-person meetings.

Taking a structured approach to experimental design, periodically reassessing, and cooperating may seem like simple steps to take, but the authors have routinely found that biosecurity efforts over the last two decades--from promoting self-governance to requiring oversight for DURC--have not taken these steps. For example, even if the DURC policies mentioned above are updated, policy revision still cannot address effective implementation, a lack of global reach, and its own list-based limitations. These hurdles in part exist because of assumptions baked into DURC, such as the idea that involvement in fundamental research governance should be severely limited to minimize the burden placed on science.

An immediate step to expand and revise these lessons is for philanthropies, governments, and others to fund a review of existing biosecurity governance experiments with the aim of determining how they are being implemented in practice. Such a review should be conducted so that the findings are able to be integrated into both policy redesign and networks of biosecurity practitioners.

In addition to the examples above, such a review should focus on industry and regions of the world that have little to no current biosecurity governance in place. The industrial and commercial development of biology represents a significant amount of biological research and innovation. Industrial organizations also have significant influence on state governance decisions, and are themselves trying out in biosecurity governance through efforts like sequence screening in the International Gene Synthesis Consortium (IGSC), which might benefit from a more experimental design. For many regions of the world without biosecurity governance, getting basic oversight capacity in place is already a major lift.

The biosecurity community should establish and strengthen shared resources to help groups wishing to establish new governance systems for their communities, such as the Analytical Approach for the Development of a National Biosafety and Biosecurity System, published by the Public Health Agency of Canada. It should also strengthen resources for cooperation and learning across regions of the world, such as the International Network on Biotechnology run by UNICRI, and the Middle East and North Africa community of practice for biological and chemical safety and security run by Gryphon Scientific.

Following through on efforts such as sharing near misses and examining the assumptions and considerations we make in designing our experiments will likely raise some concerns themselves. Publicly discussing specific instances of biosecurity concern that our governance systems do not cover can itself be an information hazard, but the *processes* of biosecurity governance may be less of a hazard to discuss. Institutions have many reasons beyond security (such as reputational and intellectual property risks) to not share information, and we encourage the exploration of options to discuss these more sensitive issues (*20*). A particularly important challenge is enabling the safe migration of useful learnings from more restricted environments (e.g., classified facilities, industrial operations) to less restricted environments (e.g., lower safety level university labs, the DIY community). Sharing an evidence base on what has and has not worked is already recognized as a necessary aspect of developing better biosecurity governance (*21*).

At present, there is no capability for *systematic* learning about the effectiveness and limitations of our current biosecurity experiments. Until we come to understand biosecurity governance as an experimental space, we will be unable to make more than sporadic movement past our current reactive governance landscape, endangering economic vitality, academic freedom, and the security of our states, people, and environment.

**References and Notes:**


1. K. M. Berger, D. DiEuliis, C. Meyer, V. Rao, *Roadmap for Biosecurity and Biodefense Policy in the United States* (Gryphon Scientific, 2017).
2. D. DiEuliis, V. Rao, E. A. Billings, C. B. Meyer, K. Berger, Biodefense Policy Analysis—A Systems-based Approach. *Health Secur.* **17**, 83–99 (2019).

3. L. Cong, F. A. Ran, D. Cox, S. Lin, R. Barretto, N. Habib, P. D. Hsu, X. Wu, W. Jiang, L. A. Marraffini, F. Zhang, Multiplex Genome Engineering Using CRISPR/Cas Systems. *Science*. **339**, 819–823 (2013).

4. C. McLeish, D. Feakes, Biosecurity and stakeholders: the rise of networks and non-state actors. *Sci. Public Policy*. **35**, 5–12 (2008).

5. L. Stampnitzky, *Disciplining terror: how experts invented "terrorism"* (Cambridge University Press, 2013).

6. NSABB (National Science Advisory Board for Biosecurity), NIH VideoCast - National Science Advisory Board for Biosecurity - January 2020 (Day 1) (2020), (available at https://videocast.nih.gov/summary.asp?live=35665&bhcp=1).

7. NIH Office and Science Policy, *Stakeholder Engagement Workshop on Implementation of the USG Policy for Institutional Oversight of Life Sciences DURC* (2017), (available at https://osp.od.nih.gov/event/stakeholder-engagement-workshop-on-implementation-of-the-usg-policy-for-institutional-oversight-of-life-sciences-durc/?instance_id=44).

8. H. Wolinsky, The FBI and biohackers: an unusual relationship. *EMBO Rep.* **17**, 793–796 (2016).

9. S. Tocchetti, S. A. Aguiton, Is an FBI Agent a DIY Biologist Like Any Other? A Cultural Analysis of a Biosecurity Risk. *Sci. Technol. Hum. Values*. **40**, 825–853 (2015).

10. R. L. Moritz, K. M. Berger, B. R. Owen, D. R. Gillum, Promoting biosecurity by professionalizing biosecurity. *Science*. **367**, 856–858 (2020).

11. F. Daviter, in *Learning in Public Policy: Analysis, Modes and Outcomes*, C. A. Dunlop, C. M. Radaelli, P. Trein, Eds. (Springer International Publishing, 2018), pp. 145–165.

12. I. van de Poel, D. C. Mehos, L. Asveld, in *New perspectives on technology in society: experimentation beyond the laboratory*, I. van de Poel, L. Asveld, D. C. Mehos, Eds. (Routledge, an imprint of the Taylor & Francis Group, Abingdon, Oxon ; New York, NY, 2018), *Emerging technologies, ethics and international affairs*.

13. P. Millett, T. Binz, S. W. Evans, T. Kuiken, K. Oye, M. J. Palmer, C. van der Vlugt, K. Yambao, S. Yu, Developing a Comprehensive, Adaptive, and International Biosafety and Biosecurity Program for Advanced Biotechnology: The iGEM Experience. *Appl. Biosaf.*, 1535676019838075 (2019).

14. B. Rappert, Why has not there been more research of concern? *Front Public Health*. **2**, 74 (2014).

15. E. M. Davidson, R. Frothingham, R. Cook-Deegan, Practical Experiences in Dual-Use Review. *Science*. **316**, 1432–1433 (2007).

16. ABSA Distance Learning Committee, (2018; https://absa.org/event/ual-use-research-of-concern-durc-roundtable-discussion/).

17. *"Looking Ahead," Presentation by Edward You, at the Workshop on BioSecurity: How synthetic biology is changing the way we look at biology and biological threats* (Synthetic Biology Project, Science, Technology, and Innovation Program, Wilson Center, Washington, D.C., 2010; https://www.synbioproject.org/events/biosecurity/).

18. A. S. Balmer, J. Calvert, C. Marris, S. Molyneux-Hodgson, E. Frow, M. Kearnes, K. Bulpin, P. Schyfter, A. MacKenzie, P. Martin, Taking roles in interdisciplinary collaborations: Reflections on working in post-ELSI spaces in the UK synthetic biology community. *Sci. Technol. Stud.* **28**, 3–25 (2015).

19. iGEM (international Genetically Engineered Machines), Gene Drive Policy (2017), (available at http://2017.igem.org/Safety/Policies).

20. NASEM (National Academies of Sciences, Engineering, and Medicine), *Dual Use Research of Concern in the Life Sciences: Current Issues and Controversies* (2017).

21. NASEM, *Governance of Dual Use Research in the Life Sciences: Advancing Global Consensus on Research Oversight: Proceedings of a Workshop* (2018).

**Acknowledgments:** The authors express their gratitude to all the participants of the *Novel Practices in Biosecurity Governance* workshop organized by Sam Weiss Evans and convened at the University of Cambridge on 11-13 July 2019 through the Biosecurity Research Initiative at St. Catharine's (BioRISC) and the Centre for the Study of Existential Risk (CSER). This document does not contain technology or technical data controlled under either U.S. International Traffic in Arms Regulation or U.S. Export Administration Regulations. The views expressed in the paper are those of the authors and not those of any institutions with which they may be affiliated; **Funding:** The workshop around which this paper materialized was supported with funding from the Hauser-Raspe Workshop Series. Sam Weiss Evans received support for drafting the article from a Schmidt Futures grant on "Ethics in the Lab" and from the Templeton World Charity Foundation; **Author contributions:** Sam Weiss Evans wrote first draft, orchestrated all edits, and designed and ran the workshop connected with this article. Jacob Beal contributed language on the industry perspective and general edits throughout. Kavita Berger rewrote substantial portions of various drafts and was heavily involved in framing discussions, as well as a panelist at the workshop. Diederik A. Bleijs edited early versions of the draft and ensured health security was incorporated. Alessia Cagnetti edited draft for flow and clarity and provided references for non-state actors. Francesca Ceroni updated references for technical literature, and edited lessons learned. Gerald L. Epstein clarified wording throughout, particularly on government-facing points. Natalia Garcia-Reyero improved references to non-state actors and technical capabilities and drafted figure. David Gillum provided revisions throughout draft, lead the ABSA working group at the workshop, was involved in framing discussions, and was a panelist at the workshop. Graeme Harkess clarified terminology and edited sections on working with less developed governance efforts. Nathan J. Hillson lead the JGI working group at the workshop and edited early drafts. Petra A.M. Hogervorst edited throughout document, particularly bringing in global perspectives. Jacob L. Jordan edited throughout, ensuring language mapped well with other NTI efforts. Geneviève Lacroix conducted extensive edits throughout the process, and was heavily involved in framing discussions. Rebecca Moritz conducted extensive edits throughout, co-lead the ABSA working group at the workshop, and was involved in framing discussions. Seán S ÓhÉigeartaigh revised the draft for clarity and flow. Megan J. Palmer was heavily involved in framing discussions, and edits throughout drafting, and co-lead iGEM working group at the workshop. Mark van Passel brought in a global health security perspective in edits, including specific clarifications on the Global Health Security Agenda, and was a panelist at the workshop; **Competing interests:** Kavita Berger is involved in running the Middle East and North Africa community of practice for biological and chemical safety and security run by Gryphon Scientific. David Gillum is incoming president of ABSA. Sam Weiss Evans and Megan J. Palmer are on the iGEM Safety and Security Committee. Mark van Passel recently stepped down as Chair of Action Package 3 for the Global Health Security Agenda. Graeme Harkess sits on the BSLG. Nathan Hillson runs the Joint Genome Institute extended screening for synthetic biology funding distribution; and **Data and materials availability:** All data is available in the main text or the supplementary materials.