Options for Management of Potentially Dangerous Information Generated by Life Sciences Research

Michael Imperiale[1] and David A. Relman[2,3,4]

[1]Department of Microbiology and Immunology, University of Michigan, Ann Arbor, MI 48103
[2]Department of Medicine, Department of Microbiology and Immunology, Stanford University School of Medicine, Stanford, CA 94305
[3]Center for International Security and Cooperation, Stanford University, Stanford, CA 94305
[4]Infectious Diseases Section, Veterans Affairs Palo Alto Health Care System, Palo Alto, CA 94304

Introduction

Communication of results of experiments is a cornerstone of the life sciences research enterprise. The sharing of not only the results but the methodology used to obtain those results is the norm. Detailed information facilitates not only the ability to reproduce the results of published experiments, but also to use those results as a foundation upon which subsequent studies can be based.

The late 20th and early 21st centuries have seen a tremendous increase in the global volume and distribution of life sciences research. This explosion of activity has been enabled by advances in technologies including laboratory procedures, bioinformatics, and computational tools, but also by the recruitment and training of many more young scientists entering these fields, as well as the rapid and wide dissemination of information via the Internet. With heightened activity has come vast improvement in technical efficiencies and capabilities. Some of the benefits that have been reaped include new biological therapeutics such as engineered monoclonal antibodies, diagnostic tests involving PCR and other molecular methods, and vaccines produced as recombinant proteins.

The so-called democratization of the life sciences has also been possible due to the low barriers to entry into the field compared to other experimental sciences: much of the equipment and supplies required to do basic, yet impactful experiments is relatively inexpensive. High school science labs are now using techniques that only highly sophisticated labs could master as recently as the 1990's. Biological data from large-scale experiments are now publically available for anyone equipped with the relevant know-how to extract new insights. Do-it-yourself biology is a growing hobby for a wide variety of individuals, many of whom belong to groups that share equipment and expertise.

The global landscape in which this life sciences research enterprise flourishes has changed concomitantly. Emerging and re-emerging infectious diseases continue to plague the world, with increasingly prominent zoonoses being attributed to population growth, movement and crowding, climate change, changes in land use, and other factors that alter the dynamics of interactions between humans and other species. Individuals and groups are becoming

increasingly emboldened to use violence against civilians in pursuit of their agendas[1]. In such a landscape, biotechnology, broadly defined, has dual uses, with the power to be enormously beneficial to human, animal, plant, and environmental health, but at the same time dangerous if misused for nefarious purposes. In the past decade or so, we have already seen examples of dual use research, motivated by public health concerns that have raised eyebrows and stimulated disparate responses with respect to information sharing. These include engineering a mouse pox virus to express IL-4 in an attempt to develop a vaccine—but instead producing a hypervirulent virus[2]; describing vulnerabilities in the US milk supply to deliberate contamination with botulinum toxin[3]; gain-of-function experiments with various avian influenza strains, some of them highly virulent, generating viruses with new mammalian transmissibility phenotypes[4]; and discovery of a new botulinum toxin in a patient that apparently could not be neutralized by existing antisera[5]. Each of these publications was handled in an ad hoc manner, involving a long, contentious debate in the case of the avian flu studies[6], to voluntary withholding of sequence information in the case of the discovery of the new botulinum toxin serotype[78].

With these considerations in mind, we propose that a re-evaluation of how life science methods and data are communicated is in order. In this paper, we focus on two main questions. First, what might trigger the desire (or need) for restrictions on dissemination of knowledge gained from life sciences experiments? Second, what are some potential mechanisms with which to manage information, and what are the important, attendant considerations?

<u>What might trigger a desire to limit the dissemination of information?</u>

---

[1] Inglesby TV, Relman DA. How likely is it that biological agents will be used deliberately to cause widespread harm? EMBO Rep 2016; 17:127-30.
[2] Jackson RJ et al., Expression of mouse interleukin-4 by a recombinant ectromelia virus suppresses cytolytic lymphocyte responses and overcomes genetic resistance to mousepox. J Virol 2001; 75:1205-10.
[3] Wein LM, Liu Y. Analyzing a bioterror attack on the food supply: the case of botulinum toxin in milk. PNAS (USA) 2005; 102:9984-9.
[4] Herfst S *et al.*, Airborne transmission of influenza A/H5N1 virus between ferrets. Science 2012; 336:1534-41. Imai M *et al.*, Experimental adaptation of an influenza H5 HA confers respiratory droplet transmission to a reassortant H5 HA/H1N1 virus in ferrets. Nature 2012; 486:420-8.
[5] Barash JR, Arnon SS. A novel strain of Clostridium botulinum that produces type B and type H botulinum toxins. J Infect Dis 2014; 209:183-91.
[6] Duprex WP, Fouchier RAM, Imperiale MJ, Lipsitch M, Relman DA. Gain-of function experiments: time for a real debate. Nature Rev Microbiol 2015; 13:58-64.
[7] Relman DA. "Inconvenient Truths" in the Pursuit of Scientific Knowledge and Public Health. J Infect Dis 2014; 209:170-2.
[8] Casadevall A, Enquist L, Imperiale MJ, Keim P, Osterholm MT, Relman DA. Redaction of sensitive data in the publication of dual use research of concern. MBio. 2013 Dec 31;5(1).

Before setting out to address this question, it is important to acknowledge that there are many steps in the research process at which the possibility of managing information needs to be discussed. First and foremost, we have argued that closer attention must be paid to whether certain risky experiments that fail to address pressing issues effectively, should be performed at all[9]. We therefore enter this discussion under the assumption that the types of information that may require scrutiny at the publication stage have arisen from experiments that were undertaken after critical review for their beneficial intent, effectiveness, and their minimization of risk. Thus, we then assume that the experiments in question (1) generated unanticipated results or results whose implications were not obvious at the outset; (2) were performed outside the purview of the usual funding review system prior to commencement; or (3) were judged to be likely to produce potentially dangerous information, but were allowed to proceed because it was deemed that the near-term benefits greatly outweighed any risks that could be reasonably anticipated at that time.

For any system of information withholding, or of information management in general, to work, there must be criteria that can be used to assess when information has "crossed a line" and is therefore subject to special actions. Attempts have been made to do this, most recently by the National Science Advisory Board for Biosecurity (NSABB) in their discussions about gain-of-function experiments involving pathogens with pandemic potential such as influenza viruses. Based on advice provided through written reports and in-person presentations to the Board, the NSABB developed a set of criteria that were proposed in its May 2016 recommendations to the US Government[10]. Specifically, NSABB noted that "research that could generate a pathogen that is: 1) highly transmissible and likely capable of wide and uncontrollable spread in human populations; and 2) highly virulent and likely to cause significant morbidity and/or mortality in humans" falls into a category that requires additional scrutiny when funding decisions are being made. However, it also noted that "once a study has been completed, it is difficult to limit the distribution of or access to the findings, particularly if the study was conducted in an open, academic environment." This followed directly from a prior report by the NSABB recommending

[9] Imperiale MJ, Casadevall A. A New Synthesis for Dual Use Research of Concern. PLoS Med 2015; 12(4): e1001813.

[10] http://osp.od.nih.gov/sites/default/files/resources/NSABB_Final_Report_Recommendations_Evaluation_Oversight_Proposed_Gain_of_Function_Research.pdf (accessed January 2, 2017)

that the harmful potential of all research falling under the definition of dual use research of concern be monitored from conception of the work through final communication[11].

One obvious limitation of this NSABB verbiage is that it applies only to gain-of-function research, and therefore only captures a fraction of research that may require additional scrutiny. Even when one expands the discussion to all 'dual use research of concern', it may not encompass the inevitable, new scenarios and developments in the future. Nevertheless, the NSABB criteria are a useful starting point. It is important to note that while much of the discussion about dual use research and its publication has focused on pathogenic infectious agents, one must not put on blinders and assume that other areas of research will not lead to a need for information control. Synthetic biology and systems biology may be two such areas. The milk supply-botulinum toxin manuscript noted above is an important case study because it did not involve wet lab research, but rather was a theoretical modeling study, and can be viewed as representative of an increasingly common type of research involving 'big data' and data mining tools. Work of this type typically arises outside of science research settings routinely subjected to biosafety and biosecurity oversight, and is typically undertaken by individuals unfamiliar with the history of biosafety guidelines.

Given this potential breadth, it is difficult to develop clear criteria that broadly define a line that ought not be crossed. That line will undoubtedly be context-dependent in many dimensions, including the area of the work, the availability of countermeasures against any potential dangers and the means to use them, and even the socio-political environment of the world at the time the work is performed. There are various factors that need to be considered, including the level of morbidity or mortality (in humans, other animals, plants), damage to supporting ecosystems, economic costs, and political implications. Nonetheless, we would like to provide a concrete example that may help frame the issues being faced.

While there are arguments on both sides of the debate as to whether one should make avian influenza strains that are highly pathogenic and transmissible by aerosol route between mammals, we think there is an analogous experiment about which there ought to be no disagreement. The question is, should someone deliberately try to isolate a mutant form of human immunodeficiency virus (HIV) that can be transmitted by aerosol route? Our answer is an unequivocal "no." First, there is no basis for concern at the present time that HIV will evolve

[11] http://osp.od.nih.gov/sites/default/files/resources/Communication_Tools%20_Dual_Use_Potential.pdf

and acquire this property through natural means, and therefore, no reasonable argument for creating such a virus in order to anticipate and study its features. (We are encouraged by a similar line of reasoning by NSABB[12]). Second, a deliberate effort to change the transmission features of HIV, if successful no matter how unlikely, could result in a pathogen whose spread through the population would be difficult to control. The risks are potentially high and the benefits nonexistent. An experiment should not be performed solely because someone finds it intellectually interesting.

Accomplishing this feat would no doubt require some serious genetic engineering: it would likely not be as straightforward as passaging a flu strain from ferret to ferret—which leads to another important consideration regarding factors that may contribute to a need to withhold, i.e., changing capabilities. Technological advances are increasing rapidly. The cost of building and running a small laboratory is low, and most of what one would need to perform experimental manipulations is readily available at home improvement stores or online marketplaces. We would argue that our ability to deal with potential harmful consequences of biological research is also lagging behind the technology. The inadequate public health response to the 2014 Ebola outbreak in West Africa illustrated what is already an extensive set of deficiencies in the state of preparedness for natural events.

Layered onto this is a consideration of the geopolitical environment that surrounds us[13]. Terrorist attacks on both military and civilian targets have been all too frequent, demonstrating that larger numbers of individuals and small groups are willing to use whatever means necessary to achieve their ends. Some have argued that the fact that none of these attacks to date have involved biological materials is an indication that terrorists deem biologics not to be a useful tool. We disagree with this assessment for at least two reasons. First, the use of biological agents to cause harm, although limited in scope and scale, has in fact occurred in recent times, including the attempt by the Rajneeshee to alter the outcome of an election in Oregon by contaminating salad bars with Salmonella bacteria, a rather unsophisticated attack, and the *B. anthracis* spore mailings right after the 9/11 attacks, which could be characterized as fairly sophisticated but narrowly targeted, to name two. Second and more importantly, we think that it is dangerous to try to put oneself inside the minds of all would-be terrorists: prior to 9/11,

---

[12] http://osp.od.nih.gov/sites/default/files/resources/NSABB_Final_Report_Recommendations_Evaluation_Oversight_Proposed_Gain_of_Function_Research.pdf (accessed January 2, 2017)

[13] See footnote 1

no one seemed to think that civilian aircraft would be used as terrorist weapons. Moreover, until fairly recently, terrorists seemed to avoid modalities that would result in their own deaths: this is no longer the case.

Next, one must consider the ease with which information travels around the globe. Until fairly recently, scientific results were published primarily in printed journals and only accessible to those who subscribed or whose institutions subscribed to those journals. Even when articles first began to be posted online, subscriptions were necessary to read the articles. With the movement towards open access ever accelerating, in 2017 anyone with a computer and internet access can read much of the scientific literature. More and more, researchers are bypassing peer review altogether and posting their studies to preprint servers, which are also freely accessible. Thus, restrictions that were imposed on access because of economic or technological reasons are far less frequent or prevalent.

In addition to these structural issues that one might consider *a priori* in devising a system for controlling information, there are events that may occur in the near or long term that could force a reactive response and a scheme for managing information that may be less thoughtful or productive. These include accidental or deliberate release of an agent from a laboratory; a bioterrorist attack; an unexpected zoonosis by a highly virulent pathogen; or development of another transformative bioengineering technology. For example, while still elusive given today's capabilities, one must wonder when predictive tools and supporting knowledge will have advanced to the point at which someone can reliably engineer additional virulence, transmissibility, altered host range, and countermeasure resistance into any existing agent, or design a novel pathogenic agent *de novo*.

For all of these reasons, we believe that it would serve us well to consider a thoughtful, deliberate plan for managing the kinds of information that will inevitably arise and pose major risks to humans, other animals, plants and their supporting ecosystems.

Options for managing dangerous information

Presently in the United States, there are two options for dealing with information generated from fundamental research: unrestricted dissemination and national security classification. These options are delineated in National Security Decision Directive (NSDD)-189, which was signed by

President Ronald Reagan in 1985 and has remained in effect ever since, with explicit affirmation by the George W. Bush administration after the 9/11 attacks[14]. NSDD-189 states that “No restrictions may be placed upon the conduct or reporting of federally-funded fundamental research that has not received national security classification, except as provided in applicable U.S. Statutes.” Most of the US departments and agencies that fund life sciences research do not support classified research, but the question that needs to be addressed is whether NSDD-189, its terminology, and its dichotomous view of the world still apply today, and whether it will continue to be relevant in the future.

Let’s first look at the definition of fundamental research in NSDD-189: "'Fundamental research' means basic and applied research in science and engineering, the results of which ordinarily are published and shared broadly within the scientific community, as distinguished from proprietary research and from industrial development, design, production, and product utilization, the results of which ordinarily are restricted for proprietary or national security reasons." We note the deliberate use of the word “ordinarily.” This means that open sharing of information is a normative value. As such, it is subject to change as norms change over time. We propose that the line between fundamental research and what the directive calls “proprietary research,” but which may also be thought of as applied research, has blurred significantly over the last several decades. Contemporary examples of such research might include fundamental studies performed at universities with the goal of commercializing the results, as has been widely pursued since passage of the Bayh-Dole Act. The use of public funds in today’s world comes with substantial social and ethical obligations and expectations—in particular, that the fruits of all such work should be reaped by the public in the form of products and goods, and that these should be monetized wherever possible. While commercialization requires disclosure in the form of patents, this does not preclude the possible need for management of the associated data. In fact, pressure to secure intellectual property rights has led to management of information derived from fundamental science. As such, there are plenty of experiments, perhaps even a majority, that are not easily categorized in one or the other bin.

The so-called Corson report, “Scientific Communication and National Security,” written in 1982 by a committee of the National Academy of Sciences and produced by the National Research Council, is helpful in providing some context here because it served as the foundation for NSDD-189. It states, “Current proponents of stricter controls advocate a strategy of security

[14] https://fas.org/irp/offdocs/nsdd/nsdd-189.htm (accessed January 2, 2017)

through secrecy. In the view of the Panel security by accomplishment may have more to offer as a general national strategy. The long-term security of the United States depends in large part on its economic, technical, scientific, and intellectual vitality, which in turn depends on the vigorous research and development effort that openness helps to nurture… Controls on scientific communication could adversely affect U.S. research institutions and could be inconsistent with both the utilitarian and philosophical values of an open society." Indeed, it has been argued that one contributor to the downfall of the USSR was a high level of secrecy and lack of sharing between entities inside the Union. However, the Corson report also clearly acknowledges that there are "gray areas" of research that lie between those that clearly deserve unrestricted dissemination and those that warrant classification, and that these 'gray areas' require special attention: "The Panel's criteria leave narrow gray areas for which, in a few instances, limited restrictions short of classification are appropriate. An example of such a gray area may be a situation, anticipated in large-scale integrated circuit work, in which on-campus research merges directly into process technology with possible military application." The criteria formulated by the Corson committee in 1982 for findings whose communication ought to be restricted (but not classified) subsumed research with dual use applications, with a short time to military application, where such applications could give short-term advantage to adversaries, and where the findings were believed to be not already held by adversaries. As one possible mechanism for information management, the committee proposed that the creators of this information voluntarily limit its dissemination.

This brings us to some practical considerations. For the purpose of the rest of this discussion, we will argue that formal national security classification presents inherent quandaries and seemingly insurmountable obstacles when considered as a control mechanism for findings that arise during the course of work born in an open, multinational, and non-governmental setting. We suggest that it is not a practical or appropriate mechanism for the kinds of information and circumstances that we are considering here, especially across a global workplace.

It could be argued that many, if not all of the examples of published studies that we raised at the beginning of this paper fall into the Corson committee's gray area. First, who will determine what information falls within that rubric? This cannot be a single person—and certainly not just the researcher—but rather will require input from a group of individuals who bring diverse perspectives to the discussion. We think this group should include respected members of the scientific, policy, and security communities, as well as other representatives of the general

public. Some thought needs to be given as to how appropriate members of the public are chosen. The group should enlist contextual and subject matter expertise as necessary. Ideally, this group would appreciate the need in some cases for taking action far in advance of the generation of the information. Either way, however, the group will need to be agile and work as rapidly as possible because if, as we expect in most cases, the information is generated by a project with recognized societal benefit, then that information must be made freely available as quickly as possible, whether it is controlled or not. We think that scientific societies, non-governmental organizations, religious and other local leaders of society, and elected members of local, regional, and national government could nominate members of this group.

A second consideration is, if information needs to be controlled, who controls it? Does it need to be in a single location or could there be multiple sites? Control by a governmental entity may not be optimal because of the risk of mistrust from the scientific and other communities, including international stakeholders. Here again there could be a role for scientific societies, including the InterAcademy Partnership[15]. Wherever the information is housed, it must be maintained with an appropriate level of physical and cybersecurity. While there is always a concern that even the most secure computer system can be hacked, we strongly believe that slowing down access to dangerous information temporarily is superior to having it freely accessible immediately.

The third consideration is, who should have access to the information? It seems that there are two categories of individuals who have a legitimate need for access to the information. The first is members of government agencies who are tasked with public, animal, agricultural, and environmental health and security, and the second is scientists who can use the information to inform future experiments aimed at enhancing the benefits of, and lessening the risks presented by the information. There needs to be a facile and transparent application and decision-making process by which individuals can request and receive access, and as is the case for discussions about whether information should be controlled, decisions about access must be expeditious. In many ways it would make sense for the same group we proposed above to make these decisions about access. We note that all members of this group will likely need a national security clearance in order to make informed decisions.

The fourth consideration is, what are mechanisms for information control that might prove practical and useful for the life sciences research workplace? While some of the policies and

[15] http://www.interacademies.net/ (accessed January 2, 2017)

practices that govern US Controlled Unclassified Information[16] may be helpful and relevant here, we suggest that further thought and consultation on this issue are needed. Some of the challenging issues that will deserve attention are standardization of practices, identification of those who control and manage this mechanism, the needs for flexibility, adaptability and review, export control, and Freedom of Information Act applicability and exemptions.

Concluding thoughts

In this paper, we have argued for a change in our thinking about how scientific information is shared. This change is needed because of a confluence of factors, including advancing technologies and technical capabilities, globalization, rapid sharing of information, and the desire of some people to cause harm to others. How to implement change is a more difficult task. As we noted, free sharing of the results of fundamental research has been the norm for a very long time. But the lines between basic and applied research have become blurred. Moreover, the very real possibility of *mis*application of information from life sciences research in the 21st century needs to be considered. Visionary leadership will be required to convince the scientific community of the need for change. We note that scientists are generally cognizant of advances in technology and how they can take advantage of those advances to pursue their research interests. As such, scientists are always agents of change. The arguments supporting a need for a change in our approach to information dissemination are strong and should not be ignored. In some ways, once a consensus is reached that change is necessary, bringing thoughtful people together to develop a usable system for managing dangerous information will be fairly straightforward. In addition, we encourage continuing education of all stakeholders about the need for change.

Finally, we have heard arguments against the need for change that we would like to address. The first is that the experiments that will generate dangerous information are so few that they can be dealt with on a case-by-case basis, or that we ought not be concerned at all. We have no doubt that these numbers will increase because of all the factors that we have described. In addition, the consequences of just one episode of deliberate misuse of information could be enormous. It would likely precipitate ad hoc, rash and poorly-conceived regulations in an attempt at information control. A second argument is that it is impossible to control information in today's world. We agree that absolute control is not feasible and therefore prefer

[16] https://www.archives.gov/cui (accessed January 2, 2017)

“management”, but as we note above, putting some barriers in the way of those who wish to do harm might either discourage them completely or at least slow their progress. Either of these outcomes can be viewed as a win for society. The third argument is that potentially dangerous information has already been published – for example the papers we noted at the beginning of this paper – and no one has yet misused that information. This is a logical fallacy. It is akin to someone in 2000 stating that since no one has ever deliberately flown a commercial jet into a large building, we don’t have to worry about it. We urge a more proactive strategy for addressing what is already a clear and pressing set of challenges.