LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-TR-740555

Workshop Summary for Maintaining Innovation and Security in Biotechnology: Lessons Learned from Nuclear, Chemical, and Informational Technologies

P. Althouse

October 25, 2017

Disclaimer

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Maintaining Innovation and Security in Biotechnology:

Lessons Learned from Nuclear, Chemical, and Informational Technologies

Workshop Summary

October 2017

CGSR

Center for Global Security Research

LAWRENCE LIVERMORE NATIONAL LABORATORY

Workshop Summary Compiled by
Joseph Johnson, Research Intern for CGSR and
Donald Prosnitz, Senior Fellow for CGSR

Executive Summary

In the fast-paced field of biotechnology where innovation has such far-reaching impacts on human health and the environment, dealing with the implications of possible illicit activities, accidents or unintended research consequences with potential detrimental societal impacts tends to remain in the background. While controls may be inevitable for the biotech industry, workshop attendees agreed that the way in which controls are implemented will play a major role in the agility and innovation of the biotechnology industry. There is little desire to slow down the pace of the gains while dealing with the security issues that arise.

As was seen from the brief examinations of the Nuclear, Chemical, and Information Technology sectors explored in this workshop, establishing a regulatory regime needs to be a partnership between the public, corporate interests, scientists, and the government. Regulation is often written to combat perceived risk rather than actual risk—the public's perceptions (occasionally even fictional portrayals) can spur regulatory efforts. This leads to the need for a thorough and continuing assessment of the risks posed by modern biotechnology. Inadequate or minimal risk assessment might expedite development in the short term but has potential negative long-term security and economic consequences. Industry and the technical community also often have a large role in setting regulatory policy, especially when well-crafted incentives are incorporated into the regulations. Such incentives might actually lead to enhanced innovation while poorly designed incentives can actually reduce safety and security.

Any regulations should be as agile and flexible as the technology they regulate and when applied to biotechnologies they will need a new framework for thinking and implementing. The new framework should consider biotechnology as a technology and not simply a science since it is an extremely complex and adaptive system. This suggests the need to invest in social and political systems, rather than exclusively in technology, as a part of maintaining not only robust, but agile regulation. Modeling and simulations to illustrate how the biotech ecosystem might respond to new policies may be a promising approach to crafting regulations.

Introduction

The Center for Global Security Research at Lawrence Livermore National Laboratory convened a workshop on August 1, 2017 on the topic of "Maintaining Innovation and Security in Biotechnology: Lessons Learned from Nuclear, Chemical, and Informational Technologies." The workshop brought together industry, government, academic and regulator representatives for this thoughtful dialog. In considering governance models, the discussions sought to consider what models have worked or not worked in the past, who or what needs to be regulated, and what metrics would measure success or failure of the regulations.

To do this, the workshop sought to gain insight and lessons learned from existing technology sectors that could be applied to the burgeoning field of biotechnology. The objectives of the workshop were to draw upon and learn from the experiences of other advancements such as

those in nuclear, chemical, and information technologies that have exhibited acute security concerns along with enormous technological promise. While life science is a mature scientific discipline, complete with rigorous peer review, quantitative methods, and historical precedent, recent advances such as synthetic biology and CRISPR/Cas9[1] have very rapidly accelerated the breadth and depth of biotechnological developments and potential applications, necessitating a fresh examination of the safety and security of biotechnology research practices. Large pharmaceutical companies have more resources than ever, while at the same time these technologies are becoming more democratized and finding their way into community labs and garages of do-it-yourself (DIY) aficionados. Indeed, this movement is advancing the life sciences, but the diaspora of innovation could become unwieldly to the point of becoming a global security and safety risk if some controls are not put in place. Against this backdrop, the conference sought to address the following questions:

- In an increasingly globalized, democratized, and open information society, what can the U.S. do to maintain its technological edge without compromising its national security?
- What can previous policies teach us on how to maintain innovation and enhance economic return while preserving a safe, secure, and socially responsible research and development community?
- Are there lessons to be learned from the regulation of previous dual-use technologies that could be applied to biotechnologies?
- In retrospect, what might have been done differently?

A keynote presentation set the scene by raising a number of questions in the context of non-pharmacological applications of modern biotechnology and the sociological issues that need to be addressed. Artificial eyes can restore sight, but also provide capabilities (such as zoom and infrared) beyond those naturally provided. How should society deal with gain of function? What are the implications of experiments that demonstrate artificially induced learning? Should society interfere in an individual's decision to acquire enhanced capabilities? Can human systems without consciousness be built to replace animal models in drug trials? What are the implications of the potential to create these "humans without brains"? The speaker encouraged the examination of the social, legal, and moral implications of a future shift from therapeutics to other applications of the new biotechnologies including gain (or loss) of function.

The panels began with a general overview of emerging biotechnology risks looking forward at the next decade, taking into account the evolution of greater public participation. The panels that followed examined the lessons learned, status, and remaining challenges in the nuclear, chemical, and information technology industries with a view toward application to the biotechnology field. The discussions were conducted under the Chatham House rule. Therefore, this summary report does not associate any comments or opinions with specific individuals.

[1] CRISPR/Cas9 is a revolutionary new genome editing tool that greatly simplifies genetic engineering in plants, animals, and humans. It has been used to make targeted genetic changes in human embryos, and can be used to construct a "gene drive" to spread genetic modifications throughout a natural population in the wild, e.g., to wipe out an entire species of mosquitos.

Panel 1: What are the challenges to maintaining innovation in biotechnology while executing safe, secure, and socially responsible policies?

Historically, the barrier to entry in commercial biotechnology was very high, requiring significant capital investment (millions to tens of millions of dollars), professionally trained (Ph.D.-level) researchers, and a lengthy research and development period before a product could be brought to market. In consequence, venture capitalists have begun to exploit the concept of "incubator" laboratories with shared facilities, and have demonstrated that significant products with real potential return can be developed with investments as small as $100,000. Often, such products include those targeted at non-drug and non-vaccine markets. Examples include producing egg whites that do not come from chickens, DNA encryption, and DNA storage. These smaller companies are able to explore the smaller markets where disruptive technologies tend to originate because the required return on investment is not nearly as high.

Biotechnology enterprises are leveraging increased use of information technology at the expense of human capital. The result is less demand for researchers at a time when the number of biology Ph.D. graduates is increasing. Some of these surplus researchers are turning to the do-it-yourself (DIY) community-based laboratories. The result is that the cost of scientific exploration has dropped while the speed of innovation has increased. In addition to the economic implications of this shift, members of the DIY biology community feel that there is a moral imperative to educate the public and "democratize" biotechnology, because these technologies are bound to have a profound impact on all our lives in the future. This community does not ask what can be done with $100,000, but with $10,000 or even $1,000. Additionally, this low barrier to entry has not resulted in less ambitious projects. Successful projects that originated using this new model include: bioprinting, bioluminescent plants, and genetic engineering. Most of these projects do not pose security (dual-use) concerns, but they do raise questions about how society will interact with biotechnology as it becomes ubiquitous. Computer technology has followed such a trajectory, where virtually everybody owns a computer, and nobody needs a license to write a new app.

Right now, advanced training is almost certainly required to convert concepts to products. However, this is changing rapidly and on an international scale. The venture capital and incubator markets, as well as community groups, have opened biotechnology to all who are interested. Curious scientists "would rather take a chance on themselves than work at [a constrictive] company like Genentech." Unfortunately, its ubiquity and utility has led to enhanced concerns over regulating DIY biotechnology with respect to general public health and safety, as well as the potential weaponization of capabilities for offensive and defensive purposes.

International agreements intended to limit biological weapons have existed for nearly 100 years (the Geneva Protocol 1925 and the Biological Weapons Convention 1975). The discovery of the Iraqi biological weapons program after the first Gulf war, Russia's confession of having developed biological weapons in the same period, the anthrax letters of 2001, and the re-emergence of smallpox and horsepox have brought another layer of urgency to reassessing existing international agreements and domestic regulations. Incidents that include the heavy use of chemicals in conventional warfare during the Iran–Iraq War of the 1980s proved that the Geneva protocol was no longer sufficient. To address this concern, export controls were

introduced for chemical weapons-relevant materials (the Australia Group) and later extended to biological materials. This list-based approach controlled the tangible flow of agents through monitoring and exclusion. However, it did not control knowledge flow, a more elusive, yet more powerful resource.

Controlling knowledge flow in biotechnology is a complex issue because, in contrast to nuclear technology, it has been developed predominately in the private sector to improve public health and combat disease. The scientific community feared that eliminating potentially harmful uses (both intentional and unintentional) would result in stunted scientific advancement and dampen innovation. These fears became a public issue when genetic engineering first became a reality in the early 1970s. Scientists convened a 1975 Asilomar Conference to establish norms of behavior and potentially avert stifling regulations on research. The conference's underlying assumption was that science is objective (value-neutral), and communal, and that scientists were motivated by the public good and could therefore be trusted to be self-governing and follow community-based standards and norms of behavior. Based upon these assumptions, the conference explicitly addressed security issues. The attitude originating there is captured today in biohackers' views that their capabilities are overestimated while their ethics are underestimated. As was the case during the Asilomar Conference, biohackers take biosafety seriously because they want to preempt concerns over safety and security.

Modern technologies (including synthetic biology and CRISPR) have made possible, in principle, the engineering of novel pathogens or toxins from scratch, which makes the select-agent list-based approach to prohibiting biological weapons problematic at best. As the repertoire of potential bioweapons expands, the U.S. is moving toward a more flexible defensive posture, often employing the same modern biological techniques to develop rapid response platforms and novel therapeutics. Examples of these defensive technologies include Advanced Development and Manufacturing of Antibody Technologies (ADAMANT) and Agile Medical Paradigm (AMP). With ADAMANT, survivors of an epidemic are sought out to quickly derive, manufacture, and disseminate antibodies. AMP allows for platforms to be modified to new threats, resulting in new vaccines within months or even weeks.

The concept of "biology as a technology" and as a "factory" emerged early in the panel discussion. This perspective has a cascading flow of repercussions, including rapid development and commercialization, and the need to focus more on fault tolerance. Biotechnology may become as ubiquitous as software development, has a growing community of both enthusiasts and aspiring professionals who have access to a field that was once exclusive to large enterprises, and has the potential to be weaponized at individual and state levels. This is an area where the United States must lead to preserve strategic stability and competitive leadership. Some take-home points from this panel:

- Dual-use concerns are important, but biotechnology is also bound to have many other profound societal impacts that demand some form of policy solution.
- Agility and rapid innovation is inherent in the new biotech business model but this diversity is largely gone from U.S. government (USG)-sponsored research. How can USG collaborate?

- The U.S. needs a thorough assessment of the risk posed by modern biotechnology.
- Treaties may limit the ability to study certain threat agents, which might interfere with developing countermeasures to those agents.

Panel 2: What can we learn from the development and evolution of nuclear technology?

In contrast to the largely internal self-regulation environment of biotechnology stands nuclear technology with narrow rather than widespread dissemination and significant national and international regulation. Early regulatory approaches aggressively and actively sought to support the peaceful uses of nuclear technology and set out mechanisms to advance and not hinder the industry.

Atoms for Peace (1953) and the Atomic Energy Act (1954) promoted civilian development, put fissionable material into the hands of the private sector, pushed the technology abroad, reduced secrecy, and ended the government's monopoly. In an effort to advance the safe use of the technology, licensing protections were put in place, and public information and basic training became available. To expand its use, financial assistance was provided to investors to build reactors at home and abroad, and the national laboratories were created to conduct research and development. The Price-Anderson Nuclear Industries Indemnity Act (1957) reduced the nuclear industry's financial risk. By the early 1960s, it was a booming industry. Orders for new nuclear plants spiked throughout the 60s and mid 70s.

In hindsight, it is easy to see a lack of the full understanding of the potential for nuclear power plants to fail and the consequences of failure. Unfortunately, risk research was not sufficiently funded and regulations were established based on the assessment of a series of maximum credible accidents. The Nuclear Regulatory Commission conducted little if any training, and safety regulations and implementation were left to the licensee. The design of containment and "defense in depth" (providing three barriers to the release of fission products—the fuel cladding, the primary coolant system envelope, and a separate, secondary containment structure) was considered to be adequate. Starting in 1965, the industry became aware that containment was not sufficient to control all accidents and regulations began to change. By the time of the 1979 Three Mile Island (TMI) accident, the public's attitude towards nuclear power had completely changed (aided by the popular media such as the movie *The China Syndrome*), and significantly enhanced regulations were enacted.

Although it is commonly believed that public opinion and new regulations killed the nuclear industry after the TMI accident, the record shows that orders for new reactors had ceased earlier. At that time, the economics had already shifted against nuclear power so, it is unclear whether the industry would have recovered even if the accident had not occurred.

The nuclear case study constitutes a progression through two related paradigms. The first paradigm, the "Learning Curve," models the "S" shape of an industry's progress. Initially, the progress is slow until the industry gains foundational knowledge, but at some point, insights and innovation cause the industry to make significant progress. The progress will be expedited by

cross-marketing and multidisciplinary synergies. As the industry matures, profitability becomes the major concern and economies of scale are reached. The supporting bureaucracy and infrastructure become technology-risk-averse in order to protect the institution. This transition squeezes out innovation and the industry stagnates.

The second paradigm, the "Talking Heads Curve," overlies the first. As the industry gains foundational knowledge, expectations start to rise. Pundits, eager to scoop the next innovation, will overestimate current industry performance in their claims, which will eventually be perceived as hype and the narrative's optimism will temper. Once the industry experiences innovation and insight, the level of performance will rise until it matches and even exceeds that in the narrative. Hence, for a period of time the pundits underestimate the technology.

The nuclear industry was a textbook example of these two models. As we saw, nuclear power was initially considered somewhat of a panacea by the Eisenhower administration, if not an exciting upgrade to the electric grid. Before the industry could match the hype, however, TMI and subsequent accidents swung the conversation on nuclear energy from potential solutions to potential dangers. Further research uncovered methods such as Probabilistic Risk Assessment (PRA) that are used to "estimate risk by computing real numbers to determine what can go wrong, how likely is it, and what are its consequences,"[2] as well as safe uses for nuclear power. Currently, the nuclear industry is considered to be in the "bureaucratization" phase of the Learning Curve and the "underestimation" phase of the Talking Heads model.

This analysis is particularly relevant to biotechnology, which is in the beginning stages of these paradigms. Biotechnology can learn from nuclear energy that "self-policing" in the early stages of industry development, although inconvenient, removes blockades to reaching the stage where innovation and insights stimulate the exponential growth that leads to acquiring sufficient foundational knowledge to advance the technology. With regulatory policies in place, pundits will not be able to obstruct progress in the field as is typical in the "overestimation" phase. Promoting industry at the expense of inadequate risk research benefits the short term, but long-term issues can have serious negative impacts.

Since public perception changes over time, either because of real or hypothesized events popularized in the media (books, movies), the biotech regulatory framework needs to be agile and understand that the technology and public perception are dynamic. Many factors influence success of innovation, including economics, public perceptions of the promise of the technology, and failure to deliver on those promises.

Panel 3: What can we learn from regulation in the chemical industry, and its impacts?

Unlike the situations of the biotechnology and nuclear industries, regulation of the chemical industry began when the industry was already mature, large, and economically significant. Like

[2] See https://www.nrc.gov/about-nrc/regulatory/risk-informed/pra.html last accessed 10-23-17.

the nuclear industry, there were concerns about the use and proliferation of chemicals in weapons of mass destruction as well as public and environmental health.

The 1899 and 1907 Hague Conventions addressed chemical warfare before these types of weapons were developed and deployed. The Hague Conventions notwithstanding, Germany introduced chemical warfare at the Second Battle of Ypres in World War I, which sparked a chemical arms race that mobilized the industry in the United States and throughout Europe. The U.S. Army's preparedness program and Chemical Corps came into existence at this time. Proponents and opponents of chemical warfare debated the future of this nascent weaponry during the discussions of the 1925 Geneva Protocol. Proponents of regulation felt that these weapons were uncontrollable and would indiscriminately lead to significant collateral damage. Opponents of the treaty felt that it would impose burdensome international regulation on the chemical industry and constrain innovation, adversely affecting the U.S. economy. The technical community united against the treaty and was able to prevent the U.S. Senate from ratifying it.

This situation had changed by the mid to later 1970s during renewed debates over U.S. ratification of the Geneva Protocol and the Chemical Weapons Convention (CWC). The proliferation of pollutants, Agent Orange, and napalm caused an enormous public outcry. By the time of the CWC negotiations, the chemical industry was eager to disassociate itself from its environmentally hostile image. Opponents of the CWC thought that it would be impossible to verify chemical stockpiles, promote and enable corporate espionage, create a burden on the industry, and make future use inevitable. However, the use of Agent Orange during the Vietnam War turned the technical community against chemical weapons, and the President was urged to halt their use. Even more significant was that the Chemical Manufacturers Association stated that the regulatory burden would be acceptable, and the industry helped shape the final language of the CWC. An economic component also factored into industry's support: it expressed concern that the U.S. would suffer a significant loss in the export market if it were not a party to the CWC. The U.S. ratified the treaty in 1993.

The most cited domestic chemical regulation is the Toxic Substance Control Act (TSCA, 1976). TSCA regulates the "manufacturing, processing, distribution, use, and disposal of commercial and industrial chemicals." TSCA emerged from environmental law and policy considerations. Still in force today, TSCA was designed to regulate chemicals products rather than applications. To expedite implementation, TSCA grandfathered approval for the use of nearly 60,000 chemicals in existence at the time of implementation—most with multiple uses. There was no requirement that these chemicals be reviewed. This regulation had some unintended negative impacts.

The two provisions resulted in an environment that worked against innovation because old chemicals could be marketed with new applications without the necessity of a TSCA review, while new chemicals required new regulatory approval. As a result, the regulation delayed release of new products at greater cost. Also, TSCA was a patchwork—it did not set out systematic deadlines and guidelines. It provided no defined limits on uses of chemicals and materials, and imposed extremely burdensome analytical conditions on the Environmental Protection Agency's attempts to regulate chemicals. These factors led to long and contentious legal battles—it took ten years and 10,000 pages to ban asbestos, and that ban was overturned

by the courts two years later. TSCA only banned 5 of the 60,000 chemicals in existence in 1976. It was amended in 2016 (the Frank R. Lautenberg Chemical Safety for the 21st Century Act) to remedy some of the faults in the original version, and to begin a review of high-priority chemicals catalogued in 1976.

Domestic safety and international security concerns intersected after 9/11 with the emergence of attention on terrorist attacks against facilities storing large quantities of toxic industrial chemicals (TICs). The requirement that industry examine “worst case” scenarios (contrasted to “maximum credible” in the nuclear arena) and make them public has led to innovation as many companies seek alternative chemicals and safer methods for achieving the same objective.

In summary, industry and the technical community played a large role in the USG regulatory stance of the chemical industry. When their positions changed, so did the USG position. External events, coupled with strong public reaction, became a primary driver resulting in a new regulatory system. For the first time the private sector sought meaningful regulation. Their motivations included improving corporate reputation, leveling the playing field, and providing a standard all could follow and were required to follow. In this way, the sector was able to preempt more rigid response. Is public perception well-enough informed to put pressure on the chemical industry in the case of fast-moving developments? This open question suggests the need to invest in social and political systems, rather than in just technology, as a part of maintaining not only robust, but agile, regulation.

Panel 4: Is regulation for information technologies a model to follow?

At first glance, information technology (IT) and nuclear power appear to be at opposite ends of the spectrum of possible regulatory approaches. The former celebrates its laissez-faire tradition of minimal regulation and the latter is portrayed as suffering stifling regulation to the point of strangulation. This perception is at least partially a function of the pervasive nature of IT applications as opposed to nuclear-limited applications. IT is more flexible, has minimal buy-in capital requirements, and is virtually ubiquitous. However, within this vast array of applications is a subset—small in proportion to the entire array, yet large compared to nuclear applications—of systems whose defects could threaten a sizeable population.

One such critical system discussed during the workshop is the control of the electric grid. After a large regional blackout (August 2003) exposed the system’s vulnerabilities, the North American Electric Corporation (NERC) and an industry group developed a set of 10 cybersecurity standards, known as the NERC–CIP (Critical Infrastructure Protection) standards. Implementing these standards was voluntary until concerns increased and they were made mandatory in 2007. This history is similar to the evolution of regulation in the nuclear industry where failures and new technical evaluations turned voluntary self-assessments into mandatory regulations. While not all problems have been eliminated (the standards do not cover all elements of the grid, e.g., the distribution systems are regulated by state, not the federal government), the multi-million-dollar fines for non-compliance focused power company executives’ attention, resulting in reliability improvements. The standards helped usher in a new cybersecurity industry along with security innovations for the electric grid.

Unfortunately, there were also unintended negative consequences. The regulation did have a leveling effect on the industry by increasing the security posture of the low performers but, arguably, decreased the posture of the best in class down to the minimum required by the new regulations. The regulations led some within the industry to move away from some existing best practices, such as maintaining “black start” compliance. Black start is the ability to restart an electric power station without relying on power from the external transmission network, and carries a heavy cost to maintain system-wide reliability. By removing black start, they were not required to stay compliant but this in turn impaired the ability of the electrical grid to rapidly restore power after a widespread blackout.

This type of regulation discourages the adoption of modern technologies because of concerns over costs resulting from new regulation—TSCA’s implementation provides an example. There, new (and presumably safer) technology brought new compliance costs that would be avoided if no changes were made. Currently, 30–70% of corporate security budgets in the power industry are dedicated to compliance, possibly resulting in the diversion of funds from new security systems.

In summary, policy does impact the reliability of the grid, but also the profit motive of the industry. Hence when security threats emerge, the regulatory approach to IT (at least one of its sectors) quickly shifts from laissez-faire to a shape and form more familiar to the nuclear and chemical industries. However, such regulation has not propagated through the IT industry and is remained confined to certain sensitive systems. Whether the broad IT industry can prevent external regulation through self-regulation remains an open question.

The public concern with the broader IT industry is related to vulnerabilities in software leading to data breaches. The incentive structure for preventing breaches does not appear to be working—there has been no decrease in the number or cost of breaches since 2008 despite regulations. The penalty imposed for breaches may actually be counterproductive since security is focused on those systems which, if compromised, results in fines, ignoring other vulnerable systems. It may be that government regulation will have little impact on incentivizing IT to solve data breach problems, but civil penalties or loss of reputation might move the industry in that direction. There is no good definition of what it means to be secure, although the National Institute of Standards continues to develop relevant standards. A better incentivized system with clear definitions and standards is needed.

Biotechnology and IT have some similarities—both systems have preexisting vulnerabilities, and breaches in both systems by negligence are common (although some expressed the opinion that biologists are a bit more cautious).

The two disciplines do have some differences—nature continually evolves the vulnerabilities to be addressed in biotech (new pathogens); the consequences of a biotech vulnerability are potentially catastrophic and irreparable compared to IT; and cyber red-teaming is viewed as improving IT security while biotech red-teaming is viewed as very dangerous and potentially a treaty violation. Finding bugs in code is celebrated in IT. Companies offer incentives to find and fix bugs. A bug in the biotech world is potentially more serious and discourages proof-of-concept experiments. However, some similarities exist: looking for vulnerabilities in the electric grid is

illegal, and tools used by individuals to find and correct other cyber vulnerabilities could be prohibited by the Wassenaar Arrangement.

Consistent with the self-regulatory nature of IT, the industry has engaged the public in the effort to find and fix vulnerabilities (bugs) by offering "bug-bounties." Unfortunately, this incentive system has a major flaw—the lack of a sound front-door mechanism for reporting bugs. Ninety-four percent of Fortune 500 companies lack this mechanism. The de facto reaction of "kill the messenger" results in discouraging an entire community of hackers from finding and reporting vulnerabilities. This is unfortunate because hackers may be best equipped for the task. The problem stems from the Computer Fraud and Abuse Act (1986), which was in part spawned by the movie *War Games*. The act indirectly makes a "crime of curiosity" and categorizes a group of hackers as "bad." In addition, the incentives for finding and reporting bugs shift the financial rewards in such a way that the defense market rewards are substantially lower than those of the offense market.

Panel 5: Achieving a practical way forward for safety and security in an evolving biotechnology industry

The objective of this panel was to guide the discussion toward some overarching conclusions and suggestions for a way forward. It was recognized that designing a biotechnology oversight system is difficult. Biotechnology regulation and innovation should be seen as a complex, adaptive system because innovation or the impact that regulatory changes might have on the entire ecosystem is hard to predict. As discussed earlier in the workshop, sometimes regulations unintentionally lead to perverse incentives. Heavy regulation is effective when there are few actors and compliance can be closely tracked. But traditional ideas of regulation may not work well in biotechnology because there are so many different interacting stakeholders, including a growing grassroots community. The biotech community also needs to further instill a culture of safety, ethics, and security, separate from regulation.

Since the consequences of policy changes can be hard to predict, agility and adaptability need to be built into the policy framework so that adjustments can be made, if needed. In the nuclear industry, the codes governing the use and release of radioactive elements are only in their second revision. In contrast, some of the chemical regulations are on their 30th revision! In biotechnology, the U.S. policy on Dual-Use Research of Concern (DURC) was always intended to be updated, but a mechanism for doing so was never designated.

Modeling and simulations to illustrate how the biotech ecosystem responds to new policies may be a promising approach. This could take the form of (1) agent-based computational models that use a risk-based approach to factor in vulnerabilities, (2) role-playing exercises (as in the Design Thinking workshop for public policy innovators at Stanford), or (3) small-scale testbeds for policy implementation (as in the annual international synthetic biology student competition iGEM, where policy rules can be adjusted yearly, allowing a design-build-test cycle).

Participants stressed the importance of transparency and buy-in from all stakeholders in developing biotechnology policy. If the process is effective, it may propagate elsewhere—U.S.

codes are often adapted by other countries around the world. But the imposition of overly restrictive regulations may drive people and innovation to other more permissive countries.

As was seen from the cases briefly explored in this workshop, establishing a regulatory regime needs to be a partnership between the public, corporate interests, scientists as political actors, and the government. In general, the group concluded that:

- Scientists had a critical role in first establishing and then changing the U.S. stance on chemical regulations. Scientists cannot ignore the potential for bad actors, either those that are careless or are intentionally conducting malicious research.
- Regulation is often written to combat perceived rather than actual risk. The media had a substantial impact on public perception and subsequent legislation and regulation as was seen after the *War Games* and *The China Syndrome* movies were released. The public perception of risk and moral outrage (e.g., following Love Canal, and the use of Agent Orange in Vietnam) dramatically changed the regulatory posture of the government. In addition, the public reaction, reputation, and market penalty for data breaches may far outweigh formal penalties.
- The regulatory regime must be flexible and adaptive to accommodate both technological and societal changes. Incentivizing industry by reducing risk estimates and accompanying regulation can have short-term benefits but long-term adverse consequences (nuclear, IT.) A thorough and continuing risk assessment is required. The regulation must be able to accommodate changing assessments of risk, and risk reduction methods.

The group noted that unintended consequences of regulation can reduce safety and security:

- Regulating the electric grid led to a focus on compliance, which distracted from true security.
- Over- (and under-) specified analytical assessments of risk creates misleading information, therefore a clear understanding of the uncertainties (and likelihood of adverse events) included in risk estimates is essential for policy-making.
- Regulation can lead to a “leveling” of safety measures to the lowest common denominator.
- Well-intentioned agreements may restrict dissemination of tools used to discover cybersecurity vulnerabilities.
- Well-meaning regulations can disincentivize security innovations because they require expensive and time-consuming approval (chemical, nuclear, and IT.)
- There is a need to invest in social and political systems, rather than in just technology, as a part of maintaining not only robust, but agile, regulation

Careful design of incentives, both corporate and individual, were considered critical:

- When the economic incentives to the chemical industry changed (potential for fines and loss of markets), the industry’s position changed.
- The current economic incentives for finding bugs favors “bad” actors.

- The current formal penalties for data breaches are nominal.

The group asked what and who should be regulated? Materials? Infrastructure? Applications? Knowledge? Data?

- Nuclear regulates materials and infrastructure. The nuclear materials-based strategy may no longer be adequate for democratized technologies.
- Chemical attempted to regulate materials. When combined with legacy approvals, this led to a damper on innovation and inadequate safety regulation.
- IT regulates applications (the grid).
- Biotechnology regulates materials (list-based select agents).

As biotechnology evolves from a laboratory science to the mainstream, it is likely to be seen more like information technology, resulting in an increase in the number and diversity of stakeholders. The emergent effects of the interaction between industry institutions, people, environments, etc. … is unpredictable. Although it is not usually characteristic of regulations, an agile and adaptive process was seen to be critical in order to continue to foster innovation while addressing safety issues. Careful consideration of lessons learned from past risk regulation can add valuable insight to the process.

Special thanks to Paris Althouse, Patrik D'haeseleer, Mona Dreicer, and Kavita Berger for their input in development of this Summary Report.

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344, LLNL-TR-740555