UNIVERSITY OF CALIFORNIA, BERKELEY
GOLDMAN SCHOOL OF PUBLIC POLICY

GSPP12-003

Taking Self-Governance Seriously: Synthetic Biology's Last, Best Chance to Improve Security

Stephen M. Maurer, Goldman School of Public Policy and Berkeley Law School, University of California.

Abstract: Synthetic biologists have vigorously debated the need for community-wide biosecurity standards for the past decade. Despite this, the US government's official response has been limited to weak and entirely voluntary Guidelines. This article describes attempts by journal editors, academic scientists, and commercial firms to organize private alternatives at the grassroots level. Private commercial standards, in particular, are significantly stronger than federal Guidelines and currently operate across more than eighty percent of the synthetic DNA industry. The paper generalizes from these examples by asking when strong private standards are both feasible and likely to produce outcomes that are comparably democratic to conventional agency regulation. It closes by describing interventions that government can use to promote and manage grassroots standards initiatives.

Keywords: Biosecurity, dual use products, self-governance, self-regulation, synthetic biology, synthetic DNA.

JEL Classifications: K29, F52.

Date: January 2012.

Contact: Stephen M. Maurer, smaurer@berkeley.edu.

I. Introduction: Synthetic Biology's Stalled Security Agenda

There is a lot to be said for thinking through a problem before you try to solve it. All the same, the situation in synthetic biology is becoming ridiculous. The modern debate over synthetic biology policy started shortly after September 11. Dozens of conferences and papers later, this work has produced almost no concrete action.[1] Indeed, the US Department of Health and Human Services ("HHS")'s main achievement during this period – developing "Guidelines" that specify what commercial gene synthesis companies should do to screen incoming customer requests for artificial DNA (HHS 2011a) – is explicitly limited to a short list of threat organisms that date from the 1960s. This backward-looking regulation has practically no chance of stopping "advanced weapons" based on DNA from other organisms.[2] Yet these are precisely the threats that synthetic biology enables and that the policy discussion was supposed to address.[3]

It would be easy to blame this failure on special interest politics. And in truth, regulators have been remarkably deferential to industry.[4] However, there is also a much deeper problem. Biosecurity – like all science and technology – is an empirical subject. But if no ideas are ever implemented, how can we learn? Practically all of the ideas currently under discussion were already being discussed in detail five years ago. *See e.g.* (Maurer et al. 2006; Garfinkel et al. 2007). Worse, policymakers have ignored much of the information that does exist. For example, gene synthesis companies have been paying human experts to screen incoming customer orders since the mid-2000s. (Maurer et al. 2009). Yet there is little or no evidence that HHS was aware of this practice when it developed its own, much less capable solution in 2009.[5] Seemingly without knowing it, HHS ended up urging industry to do less rather than more.

Plainly, the usual government channels haven't worked. Perhaps we can reform them. But in the meantime we should also look for alternatives. Over the past ten years, synthetic biologists have shown an intriguing capacity for self-regulation. At the very least, it is time to revisit the common assumption that private standards are invariably worthless.[6] To the contrary: This article argues that market forces can sometimes lead to

[1] For a detailed review of synthetic biology's biosecurity debate, *see* Maurer (2011).

[2] HHS admits that no current list is capable of "identify[ing] all or even most dangerous sequences." (HHS 2011b).

[3] HHS was originally charged with developing regulations that would apply to all threat organisms whether "Select Agents or otherwise." For details, *see* Maurer (2011) at note 227.

[4] The role of politics has been unusually naked. Indeed, the document's lead author has said that any regulation "much more onerous than what providers are currently doing…might be of some concern."(Wadman 2009). This "concern" plainly contradicts the usual instinct that government regulation is most needed in those cases where market forces prevent companies from acting responsibly.

[5] *See* Section IIC, below.

[6] This viewpoint is nicely illustrated by a 2008 *Nature* editorial which argued that the gene synthesis industry's efforts to self-govern, however "laudable," were necessarily inadequate without government

"strong self-governance" regimes in which private standards are often *more* stringent and democratically representative than government ones. Part II documents synthetic biologists' practical experience with three self-governance initiatives over the past decade. This includes, *inter alia*, a detailed discussion of how gene synthesis companies were able to create their own worldwide screening standard. Part III generalizes from this evidence to identify specific economic conditions under which qualitatively distinct, "strong self-governance" regimes are feasible. It then goes on to argue that these private standards can be just as stringent – and comparably democratic – to those issued by government agencies. Finally, it suggests policy interventions that could be used to strengthen and promote private self-governance. Part IV presents a brief conclusion.

II. Self-Governance Experiments from Synthetic Biology's First Decade

Synthetic biologists like to trace their heritage to various non-biology disciplines like electrical engineering and computer science. This inheritance is not only scientific but social. Since the middle-1990s, most decisions relating to the Internet have been made through a private institution called "The World Wide Web Consortium" or, more familiarly, "W3C." Here, a self-styled "philosopher-king" guides community discussion and – much more importantly – announces when debates have hardened into "consensus." And, remarkably, his pronouncements stick. (Berners-Lee 2000) Not surprisingly, synthetic biologists often point to W3C as a model for their own community.

But these comments are usually rhetorical. To proceed further, we need to take a closer look at how W3C achieves strong self-governance. We argue below that the key involves finding mechanisms to enforcing consensus on reluctant minorities. We then turn to three examples in which synthetic biologists tried to develop strong self-governance for themselves. On two of these occasions, the governance mechanism invoked a relatively weak mechanism – fear of criticism – to enforce consensus. We will see that these experiments produced intriguing results. In our third case the private standards were – like W3C – imposed by market forces. As already noted, the resulting standards were stronger and at least arguably more democratic than the US government's Guidelines.

A. Baseline: W3C Self-Governance.[7] Conventional self-governance typically involves a political process in which organizers try to persuade all or nearly all community members to adopt a common standard. However, this process is inherently limited so long as dissent is costless. Here, the conventional explanation is that we expect dissenters to block community standards indefinitely until organizers agree to some "lowest common denominator" outcome. But this otherwise general argument has an important loophole: In some cases, at least, dissent may *not* be costless. Probably the best-known example of

---

intervention. (Anon. 2008). Today, we know that the opposite was true: If anything, the HHS Guidelines ended up *weakening* what the private sector sought to accomplish.

[7] For a detailed an insider's account of Web governance, *see* Berners-Lee (2000).

this occurs when government threatens to intervene if the private sector fails to act.[8] However, W3C suggests a further possibility. In the New Economy, dissenters who fail to adopt an emerging standard are punished *by the market*.

W3C provides the oldest and best-known example of how such strong self-governance regimes work. The key is consumer preferences. Because the Web requires common standards, consumers know that one and only one standard will survive in the long run. And this produces the familiar "tipping" dynamic in which consumers flock to whichever standard seems most likely to win. Once this happens dissenters face a harsh choice: Drop their objections or leave the market entirely.

W3C's innovation lay in exploiting this economic "tipping dynamic" for a political purpose, *i.e.* to enforce the majority's consensus on the inevitable holdouts who would otherwise block agreement. Crucially, this makes lowest common denominator outcomes unnecessary. Once a standard reaches "critical mass" – which might or might not be an arithmetic majority – it becomes unstoppable. At this point W3C's leader can declare consensus and make it stick.

B. Organizing Biology Journals.[9] The earliest effort to organize self-governance in biology concerned so-called "experiments of concern," *i.e.* academic research that might open the door to cheaper and more effective weapons. Here, journal editors provided the natural chokepoint. After all, most academic scientists need to publish. Take away that prospect and most experiments will never be performed at all. Strikingly, journal editors had already used this logic to deny nuclear physics results to Nazi Germany at the start of World War II.[10] This inevitably raised the question of whether a similar editors' conspiracy could suppress dangerous synthetic biology results.

The initial impetus came from government. Shortly after September 11, the Department of Homeland Security and the White House announced that they were developing regulations to control the "discussion and publication" of non-classified research that might nevertheless pose a security threat. Alarmed, the American Society for Microbiology ("ASM") asked the US National Academy of Sciences to explore voluntary alternatives. This led to a high-profile August 2002 meeting between biologists and security experts. However, this initial effort failed over some scientists' objections that non-classified research should never be restricted. Five months later an impatient Bush Administration renewed the pressure by calling on the scientific community, and especially journal editors, to "come up with a process before the public demands the government do it for them." (Check-Hayden 2003).

---

[8] For an extended review of this literature, see Hoffman and Chuah (2004).

[9] Unless otherwise noted, this section follows Maurer (2011).

[10] For the definitive account, *see* Weart (1976). It is worth noting that this earlier editor's conspiracy benefited from several unique circumstances: A terrifying world-wide crisis, a relatively small community of physicists and editors, and the absence of any need to suppress data indefinitely beyond war's end. Even so, the parallel is striking.

This threat led to a highly publicized joint statement by, *inter alia*, sixteen past and present editors of leading science journals. (Journal Editors & Author's Group 2003). Crucially, the glass was only half full. While the group agreed that non-classified results should sometimes be suppressed, it said absolutely nothing about when this should happen. Instead, the decision was left to each editor's conscience. This created the obvious loophole that authors could "shop" controversial papers to different journals until some editor, somewhere agreed to publish.[11] Given enough time – and continued government pressure – this loophole might have been closed. Instead, government resolve seems to have faded as memories of September 11 grew dimmer. In the summer of 2005, *Science* wrote an editorial defending its decision to publish a controversial paper in which researchers had used synthetic DNA to "resurrect" 1918 influenza. Since *Science* had, in fact, consulted various government officials this decision did not have to be newsworthy. Instead, journal editor Donald Kennedy went out of this way announce that he would have followed his "convictions" even if the federal government had opposed publication. (Kennedy 2005) When government failed to push back, the prospects for further self-regulation more or less evaporated.

Was this outcome inevitable? Consider what would have happened if, instead of refusing the government, *Science* had announced that the paper should *not* be published or else should be published with strong restrictions. In that case, the cost of dissent for any other journal would have risen sharply. And the cost would have risen still higher to the extent that other journals publicly endorsed the decision. At this point, *Science's* judgment would have hardened into something like an enforceable standard. Of course, this is only a thought-experiment. Still, the intuition is convincing: Private standards will normally be stronger when the debate is public. Transparency matters.

B. Organizing Academic Scientists. We have already noted that many synthetic biologists are intrigued by W3C. In the early 2000s, many community members believed that the community's semi-annual Synthetic Biology 1.0 (2.0, 3.0, 4.0…) conferences could eventually evolve into a similar self-governance body.[12]

In late 2005 my University of California project polled roughly two-dozen synthetic biologists to identify ideas that could be implemented by a community-wide vote at SB2.0. This work quickly identified six initiatives that had been both widely discussed and enjoyed broad community support. In early 2006, community members debated the

[11] This loophole did not, of course, mean that the Declaration had zero value. Authors might, after all, accept modest restrictions to avoid a second round of "shopping" or to get their paper accepted at a prestigious journal like *Nature*. Still, it opened the door to an obvious "race-to-the-bottom" dynamic in which journals would compete to offer the fewest restrictions.

[12] Community leader George Church was among the most specific. Writing in 2005, he predicted that "A code of ethics and standards should emerge for biological engineering as it has done for other engineering disciplines. The community recognizes this need, but discussions are fragmentary. The next international meeting on synthetic biology (in May 2006 at the University of California, Berkeley) should make significant progress in that direction." (Church 2005).

proposals in town hall meetings webcast from the University of California and MIT. These meetings, in turn, agreed that the full community should debate and vote on the first four resolutions when SB2.0 met that May.

But this success was short-lived. Instead, conference organizers called a telephone meeting several weeks before SB2.0. The proposed vote was canceled.[13] The reasons for this decision were never announced. However, several meeting participants have told me that the reasons included concerns that the community had no agreed "constitution" for voting, that a vote could encourage hostile outside scrutiny, and/or that dissenters might "split the community" by refusing to acknowledge majority rule. Regardless, scholars have correctly seen SB2.0's failure to vote as a "failed attempt."(Parens 2005) Worse, this outcome may have deterred would-be organizers from organizing later votes. Certainly, there have been no similar initiatives from SB3.0 to SB5.0. Even so, some progress was made. When muted criticism appeared in the science press,[14] conference participants reacted by posting a draft on-line Declaration covering much the same ground as the canceled vote.[15]

As with our editors' conspiracy example, it is important to ask whether things could have turned out differently. Suppose, in particular, that the science press's criticism had been louder so that the promised SB2.0 vote was reinstated. Would the dissenters who derailed the vote in private have been similarly willing to take these positions in public? And if a vote had been taken, would the community have then *expected* more votes at SB3.0 – and every other meeting thereafter? As with our editors example, the answers are speculative. Nevertheless, it is reasonable to think that greater transparency would have made self-governance stronger.

C. Organizing Industry. In April 2008 the International Association Synthetic Biology ("IASB") held a workshop to discuss practical steps that industry could take to improve biosecurity. While IASB's membership consisted largely of German companies, several US firms also participated. The one-day meeting ended with participants agreeing to work together on multiple "work packages." The most important of these was developing an industry-wide "Code of Conduct" specifying what companies should do to screen customer orders for possible biosecurity threats. (Bernauer et al. 2008). The resulting draft Code (IASB 2009) was praised by the journal *Nature* and widely discussed in diplomatic circles. In July 2009 IASB announced that it would hold a meeting to finalize the document that November.

---

[13] Remarkably, the conference organizers were the only synthetic biologists who attended. The remaining participants were all Washington-based policy experts.

[14] For example, Aldous (2006) wrote that synthetic biologists had "reject[ed] controversial guidelines" because, *inter alia*, some community members believed that more study was needed. Service (2006) similarly wrote that the proposed options porting that options were "controversial" and "too much for synthetic biologists themselves."

[15] Second International Meeting on Synthetic Biology (2006). Though never finalized, the document prompted additional meetings and perhaps concrete work aimed at developing new technologies that companies could use to identify customer requests for dangerous DNA.

This, however, is where things got interesting. In August, two of the industry's biggest gene-makers, DNA 2.0 and Geneart, hastily assembled a competing proposal. (Check-Hayden 2009). Unlike the Code, the DNA2.0/Geneart proposal would have replaced IASB's human screeners with a pre-defined list. Because this approach facilitates automation it would have undoubtedly, in the authors' words, been "fast" and "cheap." At the same time, it would also have been less capable. This is because current lists are notoriously incomplete and will likely remain so for at least a decade. By comparison, human screeners routinely search large databases like Genbank to identify close analogs to customer orders and evaluate the functions they code for. This gives them an excellent chance of detecting threats whether or not they have seen the sequence before.[16]

This early challenge was controversial (Check-Hayden 2009) and was quietly discarded by early October. Shortly afterward, IASB members met to finalize the Code on November 3. This open meeting was attended by several non-IASB members and a reporter from *Nature* By month's end, eight companies had signed the document. Significantly, the group included two Shanghai-based companies that had not previously participated in the process. Just as importantly, DNA2.0 and Geneart joined three other large US companies to form a new alliance that called itself the International Gene Synthesis Consortium ("IGSC"). This new group was composed entirely of large companies and claimed to represent more than eighty percent of the gene synthesis industry's production capacity. By month's end, IGSC had produced a competing document which it called the "Harmonized Protocol."[17] Despite entirely new language, the Protocol appeared substantively indistinguishable from IASB's Code. *Id.*[18] This meant that most, if not all of the industry had now endorsed a human screening standard.

It would have been natural for the US (and other countries) to intervene in these developments. This could have been done by (a) endorsing a particular standard during the August-September standards war, (b) calling on the Code and Protocol groups to merge, or (c) codifying the private standards in government regulation.[19] In fact, the US did none of these things. Instead, it announced its own competing "Guidelines" at the end

[16] In practice, it is hard to see how human screening can *ever* be inferior to list-based systems. The reason is that list-based systems cost almost nothing to implement. This means that companies that use human screeners almost always run lists as well. More fundamentally, the difference between human- and list-based systems is largely an illusion – After all, lists are also compiled by humans. The real difference is that list-based systems must be compiled beforehand while "human screening" examines customer orders *as they come in*. This approach is inherently more efficient since a genuinely complete list will almost certainly include large numbers of sequences that no customer will ever ask for.

[17] The current version of the Code can be found in IGSC (2012). The original document is available from the author.

[18] The new Consortium did, however, reserve the right to amend the document in the future. *Id.*

[19] For the classic account of jawboning, see McConnell (1963). McConnell makes the important observation that jawboning is impossible in competitive industries, *i.e.* industries where members have no discretion to cut or raise prices without going out of business.

of November. (HHS 2010) Remarkably, this new document only asked gene synthesis companies to compare customer orders against genomes associated with the federal government's list of "Select Agent" organisms. If anything, this method was even weaker than the list-based proposals that DNA2.0 and Geneart had abandoned two months before. It was also much less effective than the human screening standard endorsed by the Guidelines and Code. HHS ultimately resolved this conflict by amending the final Guidelines so that they no longer address general biosecurity issues. Instead, the final document limits itself to the much narrower question of how companies should comply with the US Select Agent statute. Indeed, it expressly points out that additional biosecurity measures are both "commendable and encouraged." *Id.* This oblique reference presumably includes voluntary decisions to practice the Code or Protocol.

D. The Road Ahead. It is reasonable to ask how much life remains in any of these ventures. For now, none are particularly active:

> *Biology Journals.* Given *Science*'s repudiation of government interference (and even advice), the notion of an editors' conspiracy seems moribund. There is, however, a glimmer of hope. In late 2011, the editors of *Science* and *Nature* delayed publishing new bird flu experiments after a government panel warned that the information might pose a biosecurity threat. (Anon. 2011) The crisis ended several months later when the panel withdrew its objection on, *inter alia,* the oddly circular ground that the US government lacked any formal authority to suppress publication. (Weaver 2012) It is still not clear whether *Science* and *Nature* would have voluntarily agreed to a permanent moratorium.
>
> *Academic Scientists.* The prospects for academic self-governance remain muted. Certainly, the failure of self-governance at SB2.0 has reinforced community expectations that practical initiatives are difficult or impossible to organize. This may be why community organizers – assuming they exist – have not attempted similar votes at SB3.0, SB4.0, or SB5.0.[20] Future self-governance initiatives, if they happen at all, will almost certainly require synthetic biologists to create entirely new institutions outside the "SB" conference series.
>
> *Industry.* Based on somewhat limited evidence, most (and perhaps all) gene synthesis companies continue to practice human screening. Obvious next steps would include merging the IASB and Consortium protocols into a single standard, recruiting new member companies, developing additional, more detailed documents to further define best-practices, and/or completing a proposed open source archive where companies can share data and judgments about which DNA sequences constitute threats. There is some evidence that industry members will meet to address some or all of these issues in 2012. Absent this forward momentum, companies that have opposed strong screening in the past could

[20] The author did see members of the European Synbiosafe collaboration present their estimate of a community "consensus" at SB3.0. However, no vote was taken and claimed consensus was roundly criticized by the audience.

quietly decide to downgrade their practices. This would immediately place cost pressures on other companies to follow suit destabilizing the Code and Protocol still further.

III. Extending Strong Self-Governance[21]

Taken as a whole, our three experiments strongly suggest that W3C-style private standards can be made to work. But are more self-governance experiments feasible and/or desirable? For now, our main intuitions must come from theory.

A. Is Strong Self-Governance Feasible? It is easy to find articles that dismiss self-governance as either unenforceable or hopelessly lax. For example, Bowman and Hodge (2009) argue that self-regulation is almost always weak and sometimes little more than a public relations exercise. On the other hand, we have seen that this view fails to explain the available evidence. To the contrary: Our atomic physics and IASB examples show that self-governance can be significantly *more* stringent than formal regulation. The challenge now is to deepen our theoretical understanding of the circumstances that make strong self-governance possible.

We have argued that the key is finding mechanisms to enforce consensus in the face of dissent. Without such mechanisms, we expect even small numbers of dissenters to veto the majority's preferences and force lowest common denominator outcomes. However, the case is very different where dissent is costly. In our synthetic biology case, Big Pharma plainly did not want to do business with gene synthesis companies that failed to follow best practices. While it refused to say what those practices might be, the pressure on gene synthesis companies to adopt some unified standard was palpable. More than that, the gene synthesis companies needed Big Pharma's orders to maintain its economies of scale. This explains why companies that initially favored "cheap" and "fast" automated solutions later changed their minds. Finally, this dynamic was – like the market itself – global. Indeed, the Code's endorsement of human screening quickly spread (via the Protocol) from Europe to North America and was even adopted by two Chinese companies.

Based on our three experiments, strong self-governance seems to work best when (a) large consumers (b) who possess significant purchasing power (c) value a common standard from (d) suppliers who face significant economies of scale. By comparison, the jury is still out on whether strong self-governance can similarly be based on weaker incentives like fear of criticism. Even here, however, we expect the cost of dissent to rise steeply once it becomes public. For this reason, attempts to organize future self-governance initiatives should strive to be maximally transparent.

B. Is Strong Self-Governance Desirable? Even accepting that self-governance is feasible, we still should not pursue it unless it promises to produce democratic outcomes

[21] An extended version of the arguments presented in this section can be found in Maurer (2012).

at least as well as conventional political institutions.[22] Here, the W3C "consensus" model is more than a little disturbing. Whereas conventional democracies act on majority rule, the critical mass at which the market consensus "tips" to a single dominant standard could be much smaller. Moreover, using market forces to enforce unanimity is bound to be a rough-and-tumble process. This suggests that outcomes will often display significant randomness.

That said, we should not demand perfection. Indeed, government agencies are themselves highly imperfect, most notably in the influence they accord to large, well-organized players. The real question, then, is whether private standards would perform worse than public ones. But this presents a Catch-22 – If we never experiment with private self-governance, how can we know how well it performs? For now, IASB's Code experience provides some evidence that self-governance can produce sensible results. This is already a good reason to try further initiatives. Additionally, there are at least three theoretical arguments. First, public choice theory suggests that firms facing legal liability and/or regulatory backlash often support strong standards. If so, we should expect private self-governance outcomes to mirror the broader society.[23] Second, gene synthesis firms' willingness to let executives implement idealistic-but-expensive biosecurity standards can be seen as a form of compensation. This suggests that companies will often pursue biosecurity past the point where it ceases to be profitable.[24] Finally, bureaucracy theory suggests that government regulation tends to be disproportionately influenced by very large and/or passionate actors. Almost by definition, we expect these players to hold unusual views. By comparison, the IASB experience shows that private governance is remarkably open to small actors whose views reflect mainstream values.[25]

C. Policy Interventions: Commercial Sector. This paper has argued that stronger community self-governance is both feasible and desirable. If so, what can government do to enhance the institution? At least three comments come to mind:

Engagement. Government tends to greet private initiatives with a combination of vague praise and practical indifference. This was clearly illustrated by the fight over biosecurity screening, in which US government officials invented an entirely new standard rather

[22] The concept of democratic outcomes is, of course, fraught. For the sake of definiteness, we ask what a large representative cross-section of otherwise average citizens would decide given time to learn and consider the issues.

[23] The latter can include both formal legal liability and impacts on business from regulatory backlash. Unpublished work with Sebastian v. Engelhardt, Friedrich Schiller Universitat, Jena.

[24] This model has an interesting consequence. Suppose (a) that biosecurity costs rise linearly with the number of DNA orders filled, but (b) executive compensation rises more slowly. Then we would expect large companies to adopt lower standards than small ones. Intriguingly, this is exactly what happened in the gene synthesis industry where small European companies consistently pressed for the highest standards.

[25] This assumes that politically connected firms hold more extreme views than society as a whole. This intuition seems reasonable since (a) large firms stand to lose more from regulation than small ones, and (b) firms holding extreme views are often more passionate.

than comment on what the private sector had done. This immediately deprived the US of what would have been an opportunity to influence private standards development in Europe and China. At the same time, engaging the private standard would have forced agencies to make a clear choice and exposed them to criticism. These bureaucratic-political considerations may have made silence look like the better part of valor.

The question is how to do better. One obvious answer is to make official silence more costly than speaking out. In principle, that could be done by passing regulations that require agencies to engage private standards initiatives or else affirmatively explain why silence is desirable. But developing such regulations could take years. In the meantime, simply publicizing the need for engagement could make official silence more difficult. This could readily be done by, *inter alia*, by convening a US National Academy of Sciences panel to study the issue.

Engagement would also benefit government standards. This is because private standards contain important information. In the case of screening, for example, big companies routinely told regulators that high standards were unsustainable and would drive US employers overseas. While these claims were always debatable,[26] hearings were never likely to settle the issue. By comparison, industry's decision to adopt human screening immediately demonstrated the hollowness of these arguments. Information could also have flowed the other way. While industry clearly understood how to implement screening, it knew next to nothing about terrorism risk. Government could easily have filled this gap. All of this suggests a more general point: Traditional regulation and community self-governance should never be viewed as "either-or" alternatives. Permitting the two channels to interact will usually improve both of them.

Demanding Better. Government can accomplish a great deal simply by reminding private sector firms that self-governance is valuable and desirable. Beyond that, verbal pressure ("jawboning") could push otherwise reluctant firms to develop and implement useful standards. In keeping with our IASB example, this jawboning should extend beyond the industry to include large customers.

Transparency. Dissent is easiest when it can be done behind closed doors. This implies that self-governance is strongest when it is conducted out in the open. IASB's decision to finalize its Code in the presence of journalists provides a striking illustration of this tactic.

Supporting Market Mechanisms. Government officials typically see community self-governance in their own image, *i.e.* as a political process populated by "stakeholders" who require "engagement." In our IASB example, however, market forces were far more dominant than political ones. This suggests that policymakers should be more sensitive to economics. For example, the US has traditionally used strong export controls to restrict what American companies can sell overseas. But these same restrictions may be

[26] We have already noted that big gene synthesis companies owe their dominance to economies of scale. Cheaper labor costs abroad were never likely to erase this advantage.

counterproductive if they shelter small local companies that could not otherwise compete on price. Here, it may be better for the US to permit the emergence of a unified global market in which the same large companies – and the same private standards – operate everywhere.[27]

Establishing Partnerships. So far, we have discussed relatively simple steps that government could use to encourage self-governance. However, more ambitious collaborations are also possible. For example, we have seen that human screening is expensive. This suggests that companies can generate significant savings by sharing and re-using their screeners' Genbank searches. On the other hand, academic biologists would almost certainly find the resulting database useful. This suggests that grant authorities should be willing to provide relatively inexpensive archiving and curation resources to subsidize the effort.[28]

D. Policy Interventions: Academic Science. The most promising – and pressing – area for academic self-governance almost certainly involves experiments of concern. In analogy with our commercial screening example, government grant agencies could use their purchasing power to encourage at least three types of response:

> *Community Standards.* In principle, academic scientists could develop and adopt common standards that specify what precautions members should implement before performing experiments of concern. (Some experiments will presumably be banned regardless of precautions.) Grant agencies can promote this goal by asking applicants to specify what standards they intend to follow. This would immediately create strong pressures to create community-wide standards that anyone could use. What happened next would depend on the agency. On the one hand, government could embrace whichever standard turned out to be most popular. This would tip the community to a single universal standard, perhaps after an IASB-style standards war. Alternatively, government could decide that different experiments required different standards. In that case, the applicants might choose weak standards for some projects and strong standards for others.
>
> *Pre-Experiment Evaluations.* The experiments of concern problem is still poorly understood. For this reason, community standards could be premature. At the same time, there are obvious conflict of interest and competence reasons why grant applicants should not decide for themselves. One obvious compromise is for researchers to obtain outside, independent reviews from panels of experts.[29] Over the past ten years, there have been several attempts to create such programs.

[27] Recall that artificial DNA is a high value that can be shipped anywhere in the world by overnight delivery.

[28] US government funding of large nuclear and particle physics databases provides a compelling precedent. In both cases, data and judgments are supplied *gratis* by volunteers. Grant support is limited to storing and organizing the contributions into a coherent whole. For details see Maurer (2003).

[29] For a short survey of these programs, *see* Maurer (2011) at pp. 1414-15.

Despite this, most universities do not require reviews. Government grant authorities could change this by awarding extra points to proposals that have been vetted by independent review bodies.[30]

Post-Experiment Evaluations. Editors' willingness to withhold publication remains highly uncertain. Nevertheless, it is reasonable to think that editors may yet develop coherent uniform standards. Much will depend on community members' willingness to speak out against editors who reject government (or other editors') judgments that particular experiments should not be published.

IV. Conclusion

Synthetic biology's security discussion has produced several proposals that enjoy widespread community support. Despite this, conventional political institutions have done almost nothing to turn this consensus into action. This, in turn, has blocked the kind of trial-and-error experimentation needed to show which ideas work and which don't. Now even theory has stalled.

It doesn't have to be this way. Over the past decade, synthetic biologists have shown that self-governance can develop and implement standards that are at least as stringent as the US government's. The question now is how to expand and encourage these efforts. Here, the immediate obstacle is government. For the past ten years, government officials have routinely greeted community organizers' efforts with studied silence. Ending this silence – compelling officials to, in the jargon, "engage" private standards – would go a long way toward fixing the problem.

Scholars can help. Most observers still assume that private standards are inherently weak and/or unenforceable. In these circumstances, public officials' refusal to engage community organizers seems almost praiseworthy. But is this common assumption correct? This article has argued that – in synthetic biology, at least – private standards can be stringent, enforceable, and democratic. In the end, scholars will have to judge this claim for themselves. That said, the argument matters. Self-governance may be synthetic biology's last, best chance for improving security. Officials are much less likely to remain silent if scholars point out their error.

Acknowledgements

This work was originally presented at the SynBio and Global Health Workshop held in The Hague from February 9-10 2012. The author wishes to thank the Rathenau Institute and the Sybhel Project for supporting this work. He also thanks conference organizers

[30] An alternative, wholly private, enforcement mechanism would be for journal editors to refuse publication in cases where the experimenters failed to seek *ex ante* review. However, this would probably require community members to criticize editors who break the rule.

and participants for their comments and suggestions. Any mistakes or omissions are due to the author alone.

References

Aldous P (2006) “Synthetic Biologists Reject Controversial Guidelines. *New Scientist* May 23.

Anon. (2011) Scientists asked not to reveal bird flu details. *Wall Street Journal* (Dec. 22).

Anon. (2008) Editorial: Pathways to security. *Nature* 455:432 (2008).

Bernauer H and J Christopher, W Deininger, M Fischer, P Habermeier, K Heumann, S Maurer, H Schwer, P Stähler, and T Wagner (2008) Workshop Report: Technical solutions for biosecurity in synthetic biology.” http://www.ia-sb.eu/tasks/sites/synthetic-biology/assets/File/pdf/iasb_report_biosecurity_syntheticbiology.pdf. Accessed April 6 2012.

Berners-Lee T (2000) *Weaving the web: The original design and ultimate destiny of the world wide web.* Harper Collins, New York.

Bowman D and G Hodge (2009) Counting on codes: An examination of transnational codes as a regulatory governance mechanism for nanotechnology. *Regulation and Governance* 3:145-54.

Check-Hayden E (2009) Keeping genes out of terrorists’ hands. *Nature* 461:22.

Check-Hayden, E (2003) US officials urge biologists to vet publications for bioterror risk. *Nature* 421: 197.

Church, G (2005). Let Us Go Forth and Safely Multiply. *Nature* 438:423.

Garfinkel M, D Endy, GL Epstein and RM Friedman (2007) Synthetic genomics: Options for governance. http://www.synbiosafe.eu/uploads///pdf/Synthetic%20Genomics%20Options%20for%20Governance.pdf. Accessed April 6 2012.

Hoffmann R and Chuah SH (2004) Industry self-regulation: A game-theoretic typology of strategic voluntary compliance. *Intern. J. Econ. Bus.* 11(1): 91–106.

IASB. 2009. Code of conduct for best practices in synthetic biology. http://www.ia-sb.eu/go/synthetic-biology/activities/press-area/press-information/code-of-conduct-for-best-practices-in-gene-synthesis/. Accessed April 6, 2012.

IGSC. 2012. Harmonized screening protocol: Gene sequence & customer screening to promote biosecruity. http://www.genesynthesisconsortium.org/wp-

content/uploads/2012/02/IGSC-Harmonized-Screening-Protocol1.pdf. Accessed April 6, 2012.

Journal Editors & Authors Group (2003) Statement on scientific publication and security. *Science* 299:1149 (2003).

Kennedy, D (2005) Editorial: Better never than late. *Science* 310:195.

Maurer S (2012). *Regulation without government: European biotech, private anti-terrorism standards, and the idea of strong self-governance*, (Nomos: Baaden-Baaden).

Maurer S (2011) End of the beginning or beginning of the end? Synthetic biology's stalled security agenda and the prospects for restarting it. *Valparaiso Univ. L. Rev.* 45(4): 1387-1446.

Maurer, S, M Fischer, H Schwer, C Stähler, and P Stähler (2009) Making commercial biology safer: What the gene synthesis industry has learned about screening customers and orders. http://gspp.berkeley.edu/iths/Maurer_IASB_Screening.pdf. Accessed April 6, 2012.

Maurer S, K Lucas and S Terrell (2006) From understanding to action: Community-based options for increasing safety and security in synthetic biology. http://gspp.berkeley.edu/iths/UC%20White%20Paper.pdf. Accessed April 6, 2012.

Maurer, S. (2003) New institutions for doing science: From databases to open source biology." http://www.merit.unimaas.nl/epip/papers/maurer_paper.pdf.

McConnell G (1963) *Steel and the Presidency – 1962.* WW Norton, New York.

Parens, E, J Johnston and J Moses (2009) Ethical issues in synthetic biology: An overview of the debates. Woodrow Wilson International Center for Scholars, Washington DC. http://www.synbioproject.org/library/publications/archive/synbio3/. Accessed April 6, 2012.

Second International Meeting on Synthetic Biology (2006), Declaration. http://syntheticbiology.org/SB2Declaration.html. Accessed April 6, 2012.

Service RF (2006) "Synthetic biologists debate policing themselves," *Science* 312:1116.

HHS (2010a) "Screening Framework Guidance for Synthetic Double-Stranded DNA Providers," *Federal Register* 75(197): 62820-62832.

HHS (2010b), Frequently Asked Questions: Screening Framework Guidance for Providers of Synthetic Double-Stranded DNA. http://www.phe.gov/Preparedness/legal/guidance/syndna/Documents/synbio-faq.pdf. Accessed April 6 2012.

NIH (2011). NIH News: Press statement on the NSABB review of H5N1 research, (Dec. 20). http://www.nih.gov/news/health/dec2011/od-20.htm**.** Accessed April 5 2012.

HHS (2009) Screening framework guidance for providers of synthetic double-stranded DNA. *Fed. Reg.* 74(227): 62319-62327.

Wadman M (2009) US drafts guidelines to screen genes. *NatureNews* (Dec. 4) http://www.nature.com/news/2009/091204/full-news.2009.1117.html.

Weart SR (1976) "Scientists with a secret," *Physics Today* 29 (2): 29-30.

Weaver, C. (2012) Panel backs publishing bird-flu research *Wall Street Journal* (March 30).